


PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 7 : H04L 12/64, 12/56		A1	(11) International Publication Number: WO 00/54469
			(43) International Publication Date: 14 September 2000 (14.09.00)
(21) International Application Number: PCT/US00/06481			(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 10 March 2000 (10.03.00)			
(30) Priority Data: 60/123,798 11 March 1999 (11.03.99) US 09/405,993 27 September 1999 (27.09.99) US			
(71) Applicant: CISCO TECHNOLOGY, INC. [US/US]; 170 West Tasman Drive, San Jose, CA 95134 (US).			
(72) Inventors: KAO, Ruey; 13502 Myren Drive, Saratoga, CA 95070 (US). SUE, Cherng-Ren; 6066 Prince Drive, San Jose, CA 95129 (US). YILDIRIM, Cem; 6405 Samar Drive, San Jose, CA 95119 (US).			
(74) Agents: VINCENT, Lester, J. et al.; Blakely, Sokoloff, Taylor & Zafman, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025-1026 (US).			

Published*With international search report.**Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.*

(54) Title: MULTIMEDIA AND MULTIRATE SWITCHING METHOD AND APPARATUS



(57) Abstract

A switch apparatus and method according to the invention implements a three stage switching process. Various types of media streams presented to the switch apparatus by the broadband and narrowband connections are adapted for switching by being converted to ATM cells and enqueued in corresponding virtual circuit (VC) queues. ATM cell switching is performed among the different cards based on the quality of service requirement for each virtual circuit. The switched ATM cells are then converted to the outgoing media types and outputted to the necessary broadband and narrowband connections. The switch apparatus and method is further adapted to perform rate shaping and traffic management so as to guarantee the quality of service for various media types (voice, video, data) and also minimize the traffic loss due to rate mismatch between narrowband and broadband connections during the burst period. By virtue of this implementation, the switch apparatus and method of the present invention can perform any-to-any media type switching.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

-1-

MULTIMEDIA AND MULTIRATE SWITCHING METHOD AND APPARATUS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to carrier class switches, and in particular, to a method and apparatus for providing multimedia and multirate switching in an integrated carrier-class switching platform.

2. Description of the Related Art

Most conventional switches are adapted for switching traffic of a specific media such as voice switches (e.g. 4ESS, DMS200) and data switches (e.g. ATM/Frame Relay), or for switching traffic of a specific speed such as backbone switches (switching traffic among broadband interfaces, e.g. OC-12) and access switches (switching traffic among narrowband interfaces, e.g. DSO).

In particular, a backbone switch typically includes a switching fabric that switches between a certain number (e.g. 8 or 16) of high speed ports such as OC-12 ports. In contrast, low speed traffic is generally switched between low speed ports in an access switch having, for example, a shared bus architecture. For such low speed traffic to access one of the high-speed lines coupled to the switching fabric (and vice-versa), access switches generally include an uplink module to convert the low-speed traffic to high-speed traffic, and a separate OC-12 line is needed to transmit the converted low-speed traffic to the high-speed switching fabric. Such methods typically require one high speed port to be dedicated for each

-2-

low-speed uplink connection, thus wasting the available bandwidth on the high-speed port.

Moreover, traffic can typically be switched within one type of network. For example, ATM traffic can only be switched among ATM networks, IP traffic can only be switched among IP networks, etc.

Accordingly, there remains a need in the art for an integrated switching apparatus that provides for switching among both low speed ports (e.g. NxDSO) and high speed ports (e.g. OC-3). Moreover, such a switching apparatus should be able to minimize the traffic loss due to rate mismatch during the burst period.

There further remains a need in the art for a switch apparatus that provides for switching among different types of multimedia streams. Such a switching apparatus should also be able to guarantee the quality of service for the different media types, e.g. voice, video, data.

SUMMARY OF THE INVENTION

Accordingly, an object of the present invention is to provide a switching apparatus and methodology that permits traffic of any speed to be switched between ports of a single device.

Another object of the present invention is to provide a switching apparatus and methodology that permits traffic of any media to be switched in a single device.

Another object of the present invention is to provide a switching apparatus and methodology that permits traffic of any network to be switched in a single device.

-3-

Another object of the present invention is to provide a switching apparatus and methodology that minimizes traffic loss due to rate mismatch between narrowband and broadband connections during the burst period.

Another object of the present invention is to provide a switching apparatus and methodology that guarantees quality of service for flows of various media types.

To achieve these objects and others, the switch apparatus and method according to the invention implements a three stage switching process. Various types of media streams presented to the switch apparatus via broadband and narrowband flows (e.g. voice/fax call, video session, packet flow between source and destination ports, etc.) are converted to ATM cells and enqueued in corresponding virtual circuit (VC) queues. ATM cell switching is performed among the different cards based on the quality of service required for each virtual circuit. The switched ATM cells are then converted to the outgoing media types and outputted to the necessary broadband and narrowband flows. The switch apparatus and method are further adapted to perform rate shaping and traffic management so as to guarantee the quality of service for various media types (voice, video, data) and also minimize the traffic loss due to rate mismatch between narrowband and broadband connections during the burst period. By virtue of this implementation, the switch apparatus and method of the present invention can perform any-to-any media type switching.

-4-

BRIEF DESCRIPTION OF THE DRAWINGS

These and other objects and advantages of the present invention will become apparent to those skilled in the art after considering the following detailed specification, together with the accompanying drawings wherein:

FIG. 1 illustrates an implementation of a switch apparatus according to the present invention;

FIG. 2 is a block diagram of a switch apparatus according to the present invention;

FIG. 3 is a block diagram further illustrating an example of a local switch module that can be included in a switch apparatus of the present invention such as that shown in FIG. 2;

FIG. 4 is a block diagram further illustrating an example of an ingress module that can be included in a local switch module such as that shown in FIG. 3;

FIG. 5 is a block diagram further illustrating an example of an egress module that can be included in a local switch module such as that shown in FIG. 3;

FIG. 6 is a block diagram further illustrating an example of a narrowband service card that can be included in a switch apparatus of the present invention such as that shown in FIG. 2;

FIG. 7 is a block diagram further illustrating an example of a voice/fax controller module that can be included in a narrowband service card such as that shown in FIG. 6; and

FIG. 8 is a block diagram further illustrating an example of a DSP controller module that can be included in a voice/fax controller module such as that shown in FIG. 7.

-5-

DETAILED DESCRIPTION OF THE INVENTION

FIG. 1 illustrates an exemplary implementation of the present invention. As illustrated, a switch apparatus 102 according to the present invention provides, within a single device, the capability to switch flows of any media among LEC/CLECs, PSTNs, and ATM and IP networks.

An example of a switch apparatus 102 according to the present invention such as that shown in FIG. 1 is further illustrated in FIG. 2.

As shown in FIG. 2, switch apparatus 102 includes an ATM cell switching fabric 208 that switches ATM cell traffic between switch ports 214-1...214-N. Coupled to one of the switch ports is a broadband service card (BSC1) 210 for interfacing with a plurality of broadband connections. Coupled to a second one of the switch ports is a switch control card (SCC) 212. Coupled to another of the switch ports is a local switch module 206. Further coupled to the local switch module 206 is a plurality of narrowband line cards (NMC1, NMC2) 202, 204 for interfacing with a plurality of narrowband connections.

ATM cell switching fabric 208 is, for example, an ATLANTA chipset switch fabric having an 8x8 array of switch elements such as LUC4AS01 ATM Switch Elements made by Lucent Technologies of Allentown, PA. Such a switch fabric switches ATM traffic between eight switch ports 214-1...214-8 (only two such ports are shown in FIG. 2 for clarity). Switch ports 214 are preferably OC-12 or equivalent ports. An implementation of such an ATM cell switching fabric is described in Lucent Technologies Product Brief, "ATLANTA ATM Switch Core Chip Set," March 1997, the contents of which are incorporated herein by reference.

-6-

BSC1 210 provides an interface between one switch port 214-2 of the ATM cell switching fabric 208 and one or more broadband connections such as T3/E3, OC-3, and OC-12 lines and/or ports. Although only one broadband service card is shown, it should be apparent that there may be several.

SCC 212 contains functionality for establishing, routing, and managing virtual circuit connections between the ports of the switch apparatus. An example of an apparatus and method that can be used to implement such functionality is described in co-pending U.S. Provisional Appln. No. 60/_____, Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference.

Local switch module 206 provides an interface between one switch port 214-1 of the ATM cell switching fabric 208 and one or more narrowband line cards NMC1 202, NMC2 204. The switch port 214-1 of the ATM cell switching fabric 208 that is coupled to local switch module 206 is configured to control, for example, 16 multiPHY devices on the physical layer side. This can be implemented using, for example, a LUC4AU01 ATM Layer UNI Manager (ALM) from Lucent Technologies (not shown). Transfers of ATM cells between ATM cell switching fabric 208 (via ALM) and local switch module 206 are preferably performed via a 16-bit UTOPIA II interface (not shown). The local switch module 206 thus allows all the narrowband connections to share the bandwidth of one broadband connection. This improves the prior art solution of separately adapting one broadband connection for each narrowband interface.

-7-

NMC1 202 and NMC2 204 each provide an interface between local switching module 206 and one or more narrowband connections such as NxDSO, NxT1/E1, Ethernet, ISDN lines and/or ports. Although two narrowband interface cards are shown, it should be apparent that there may be one or several.

Generally, the switch apparatus illustrated in FIG. 2 implements a three stage switching process. Various types of media streams presented to the switch apparatus by the broadband and narrowband flows are adapted for switching between connected networks by being converted to ATM cells and enqueued in corresponding virtual circuit (VC) queues. ATM cell switching is performed among the different cards based on the quality of service required for each virtual circuit. The switched ATM cells are then converted to the outgoing media types and outputted to the necessary broadband and narrowband connections. The switch apparatus is adapted to perform rate shaping and traffic management so as to guarantee the quality of service for various media types (voice, video, data) and also minimize the traffic loss due to rate mismatch between narrowband and broadband connections during the burst period. By virtue of this implementation, the switch apparatus of the present invention can perform any-to-any media type switching as listed in Table 1.

Table 1 Switching Matrix

Input\Output	Voice/Fax/Video	Frame (FR, PPP)	ATM Cell	LAN
Voice/Fax/Video	Voice Switching	VoFR/VoIP	VoATM VoIP	VoIP
Frame (FR, PPP)	VoFR/VoIP	Frame Switching	FR/ATM Interworking	Encapsulation e.g. RFC 1490
ATM Cell	VoATM/VoIP	FR/ATM Interworking	Cell Switching	Encapsulation e.g. RFC 1483
LAN	VoIP	Encapsulation e.g. RFC 1490	Encapsulation e.g. RFC 1483	LAN Switching

FIG. 3 further illustrates a local switch module 206 such as that included in the switch apparatus 102 shown in FIG. 2. As shown in FIG. 3, local switch module 206 includes shared buses 306-A and 306-B that are coupled to NMCs for communicating ATM cells. Local switch module 206 further includes an ingress module 302 and an egress module 304 that are responsible for interfacing ATM cells between the low-speed NMCs and the high-speed ATM cell switching fabric. Preferably, the ingress and egress modules are separately embodied as FPGAs.

Shared buses 306-A and 306-B are preferably each a Cubit-Pro CellBus from TranSwitch Corp. of Shelton, CT. As will be explained in more detail below, the shared buses can be configured for load sharing mode, wherein both buses are active at the same time, or they can be configured for redundancy mode, wherein only one of the buses is active.

FIG. 4 further illustrates an ingress module 302 such as that included in the local switch module 206 shown in FIG. 3.

-9-

As shown, ingress module 302 provides interfaces between the ATM cell switching fabric and three PHY devices. The first interface is with the 16 bit ingress UTOPIA II bus of ALM 402 (such as a LUC4AU01 ATM Layer UNI Manager from Lucent Technologies, for example) of ATM cell switching fabric 208. The ingress module is slave to ALM 402 and responds to 16 PHY addresses coming from the ALM. It runs with a 50 MHz clock 450 generated by the ALM and operates in the cell by cell mode. Cells are 27 words long.

The second interface is with the 8 bit UTOPIA bus in which the ingress module is master to the SAR chip 422 (such as an L64364 ATMizer II+ from LSI, for example) TX Utopia. It runs with the 50 MHz clock generated by the ALM and operates in the cell by cell mode. Cells are 53 bytes long. The SAR (ATMizer) chip 422 is used for segmentation and reassembly of AAL5 frames (used for, e.g., intercard communication, SVC, etc.)

The third and fourth interfaces are with the 16 bit UTOPIA-like inlet buses of Cubit chips 424-A and 424-B (such as TXC-05802 Cubit-Pro Cell Bus switches from TranSwitch Corp. of Shelton, CT) are slaves to the ingress module 302. These buses run with a 40 MHz clock 452 and operate in cell by cell mode. Cells are 28 words long including a 1-word routing tag appended in front of them.

As shown, further included is a CPU interface 426 for allowing configuring and error handling of ingress module 302 by an external CPU.

In this example of the invention, the ingress module 302 emulates 16 PHY devices associated with one port of the ATM cell switching fabric

-10-

208. According to the setup of these 16 UTOPIA addresses, incoming cells are mapped to the three different data paths. Corresponding to each data path, the ingress module 302 maintains three external synchronous FIFOs 404, 406, 408 for cell storage. Each FIFO can store 8 cells. They also provide clock synchronization for the data. The ingress module 302 generates the control signals for these FIFOs.

In general, cells are received from ALM 402 and stored in the proper FIFOs depending on their UTOPIA addresses and the running mode of ingress module 302. The ingress module 302 can be configured for load sharing, in which both Cubit chips 424-A and 424-B are active, or redundancy, where only one of the Cubit chips is set as the active Cubit chip. If one or more of the FIFOs 404, 406, 408 are almost full, the ingress module will not allow the ALM chip to send more cells to those specific FIFOs. This will be achieved by deasserting cell space available signals when responding to ALM's Utopia addresses.

Ingress module 302 further periodically polls SAR (ATMizer) chip 422 and the Cubit chips 424-A and 424-B to see if they can receive a cell. If they can, one cell will be read from the corresponding FIFO and sent to the proper PHY chip. This will be done concurrently for each chip.

For cells going to SAR (ATMizer) chip 422, 27 word cells are converted into 53 byte cells by removing the UDF2 byte of the cells. For cells going to Cubit chips 424-A, 424-B, the first word of the cell coming from ALM 402 is sampled and a Routing Tag is generated out of the first word. The last 4 bits of the Tag are generated from the first 12 bits of the generated Tag with CRC-4 calculation over it.

-11-

As shown in FIG. 4, ingress module 302 includes a data path module 410, an ALM TX Utopia controller module 418, an ATZ FIFO controller module 412, Cubit FIFO controller modules 414 and 416, and a reset and CPU interface module 420.

Data path module 410 receives cells from the FIFOs and performs the word to byte conversion for the SAR (ATMizer) chip 422 and the TAG generation and insertion for the Cubit chips 424-A, 424-B. Control signals for the muxes come from the respective FIFO controller modules 412, 414, 416.

The ALM TX Utopia controller module 418 interfaces with the ALM's ingress UTOPIA bus control signals portion. It also receives cell space available signals from the FIFO control modules. During normal operation, ALM 402 continuously puts the addresses of the 16 PHY devices on the Utopia address bus. ALM TX Utopia controller 418 responds to this polling by putting the cell space available signals coming from the FIFO control modules for each PHY device. These signals are first synchronized to the ALM's 50 MHz clock 450. If ALM 402 starts to send a cell to one of these PHY devices, ALM TX Utopia controller module 418 will first check an SOC alignment. If SOC is not aligned with the cell, it will be ignored. If the cell is normal, it will be written into the FIFO corresponding to the UTOPIA address and the existing setup, as described below.

If load sharing is enabled, cells with a PHY device number 0-7 will be written into Cubit Pro A FIFO 406, and cells with a PHY device number 9-15 will be written into Cubit Pro B FIFO 408. If load sharing is disabled and if Cubit Pro 424-A is active, cells with a PHY device number 0-7 and 9-15 will be written into Cubit Pro A FIFO 406. If load sharing is disabled

-12-

and if Cubit Pro 424-B is active, cells with a PHY device number 0-7 and 9-15 will be written into Cubit Pro B FIFO 408. Cells with PHY device number 8 will always be written into ATZ FIFO 404. After the cell is written into the proper FIFO, one increment cell counter signal will be sent to the proper FIFO controller module.

ATZ FIFO controller module 412 runs at the ALM's 50 MHz clock and keeps the cell counter for the ATZ FIFO 404 and provides interfaces to the ALM Utopia bus. If the SAR (ATMizer) chip 422 has cell space available and the corresponding FIFO cell counter indicates that there is a cell available in FIFO 404, this module starts to read the cell from FIFO 404. During this process it will also generate the control signals for the data path module 410 for muncing and UDF2 byte removing. ATZ FIFO controller module 412 assumes that cells were written into the FIFO with proper SOC alignment (this is the responsibility of ALM TX Utopia controller module 418). But if something goes wrong and a cell comes out of the FIFO with a wrong alignment, this module will generate an SOC Error signal to the Reset Module 420. It will also generate a cell space available signal to the ALM TX Utopia module 418 when appropriate.

Cubit FIFO controller modules 414 and 416 respectively manage Cubit Pro chips 424-A and 424-B running at 40 MHz. They keep the cell counters for the FIFOs and provide interfaces to the Cubit 424-A and Cubit 424-B 16-bit cell inlet buses.

If the Cubit chips 424-A, 424-B have cell space available and the corresponding FIFO cell counter indicates that there is a cell available in one of FIFOs 406, 408, these modules start to read the cell from the corresponding FIFO. During this process they will also generate the control

-13-

signals for the data path module 410 for muxing and tag insertion. The Cubit FIFO controller modules 414, 416 assume that cells were written into the corresponding FIFO 406, 408 with proper SOC alignment (this is the responsibility of ALM TX Utopia controller module 418). But if something goes wrong and a cell comes out of the FIFO with a wrong alignment, these modules will generate SOC Error signals to Reset module 420.

Reset and CPU interface module 420 generates a reset sequence for the following conditions:

(1) Missing SOC signal coming from ATMizer or Cubit FIFOs for cells going from the FIFOs to the PHY devices. This is an error condition on FIFO control, not a missing SOC on the ALM Utopia side to FIFO, in which case the cell is discarded.

(2) ATZ or Cubit chip FIFO overrun condition detected. Normally, the FIFOs can not overflow. When they are almost full, the ALM chip will not be allowed to send more cells using the Utopia control signals.

(3) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

During the reset sequence, the Utopia interfaces are blocked by not allowing any cell operation. Internal logic is reset and the external FIFOs are flushed by generating proper reset signals for them.

Reset and CPU interface module 420 also provides the following pins for access by an external CPU (via CPU interface 426):

(1) Main_reset_ is an asynchronous active low input signal. It has to be asserted for a minimum period of 1 ms. This will initiate a

-14-

reset sequence, which resets the chip, indirectly resets the FIFOs and stops the Utopia interfaces. After main reset is de-asserted, it will take 32 clocks to begin operational mode. At the end of the reset sequence, pin "Err_stat" will be set as an acknowledgment, which signal needs to be cleared using the pin "Clear err cond".

(2) "Err stat" is an asynchronous active high output signal. It will be latched for the following conditions.

(a) Missing SOC signal for cells coming from ATMizer or Cubit chip FIFOs to the PHY devices. This is an error condition on FIFO control. This is not a missing SOC on cells coming from the ALM Utopia side to FIFO, in which case the cells will be discarded and the hardware will realign itself to the in-coming cells by searching for cells with a proper SOC indicator.

(b) FIFO overrun condition detected. Normally the FIFOs can not overflow. When they are almost full, the ALM chip will not be allowed to send more cells to them.

(c) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

When these conditions (a), (b) and (c) are detected, an auto initialization reset sequence is initiated. This will reset the chip and the FIFOs. This will also hold the Utopia interfaces in the blocking state, which will not allow any cell transfer until the reset sequence is completed.

This signal can be cleared with the signal "clear_err_cond".

(3) "Clear_err_cond" is an asynchronous input signal. It will clear the "Err stat" output pin.

-15-

(4) "Share" signal is an asynchronous active high input signal.

0 = Sharing is disabled (default)

1 = Sharing is enabled. In this case, UTOPIA addresses indicating PHY device numbers 0-7 will cause corresponding cells to be directed to Cubit pro 424-A, and UTOPIA addresses indicating PHY device numbers 9-15 will cause cells to be directed to Cubit pro 424-B

(5) "Act_cbt" is an asynchronous input signal. It shows which Cubit pro is the active chip. If sharing is disabled it will have the following functionality:

0 = Cubit pro 424-A is active: UTOPIA addresses indicating PHY device number 0-7 and 9-15 will cause cells to be directed to Cubit pro 424-A

1 = Cubit pro 424-B is active: UTOPIA addresses indicating PHY device number 0-7 and 9-15 will cause cells to be directed to Cubit pro 424-B.

This signal is ignored if sharing is enabled.

After power up, reset and CPU interface module 420 comes up in the reset mode (since Main reset is an active low signal) and stays in the reset mode until Main reset is de-asserted. Before this signal is de-asserted, the ALM, Cubit and SAR (ATMizer) chips need to be initialized to avoid any initial cell loss.

FIG. 5 further illustrates an egress module 304 such as that included in the local switch module 206 shown in FIG. 3.

The egress module 304 provides interfaces between the ATM cell switching fabric 208 and three PHY devices. The first interface is with the

-16-

16 bit egress UTOPIA II bus of the ALM 502 (such as a LUC4AU01 ATM Layer UNI Manager from Lucent Technologies, for example) of the ATM cell switching fabric 208. The egress module 304 is slave to the ALM and responds to single PHY address coming from the ALM 502. It runs with a 50 MHz clock 550 generated by the ALM and operates in the cell by cell mode. Cells are 27 words long.

The second interface with the 8 bit Utopia bus in which the egress module is the master to the SAR chip 526 (such as an L64364 ATMizer II+ from LSI, for example) RX Utopia. It runs with the same 50 MHz clock 550 generated by the ALM and operates in the cell by cell mode. Cells are 53 bytes long. The SAR (ATMizer) chip 526 is used for segmentation and reassembly of AAL5 frames (used for, e.g., intercard communication, SVC, etc.).

The third and fourth interfaces are with the 16 bit Utopia-like Cubit-pro outlet buses of Cubit chips 524-A and 524-B (such as TXC-05802 Cubit-Pro Cell Bus switches from TranSwitch Corp. of Shelton, CT) are slaves to the egress module 304. These buses run with a 40 MHz clock and operate in cell by cell mode. Cells are 27 words long with no Routing tag.

As shown, further provided is a CPU interface 516 for allowing an external CPU access for configuring and error handling of egress module 304.

In this example of the invention, egress module 304 emulates a single PHY device associated with one port of the ATM cell switching fabric 208. Accordingly, the three PHY devices are mapped to a single PHY address.

-17-

Egress module 304 uses three external synchronous FIFOs 518, 520, 522 for cell storage corresponding to each of the three data paths. The ATZ FIFO 518 can store 8 cells. Cubit FIFOs 520, 522 can store 8 or 144 cells depending on the configuration. The FIFOs also provide clock synchronization for the data. The egress module 304 generates the control signals for the FIFOs.

In general, cells are received from ATMizer or Cubit chips and stored in the proper FIFOs. If one or more of the FIFOs 518, 520, 522 are almost full, the egress module 304 will not allow the PHY chips to send more cells to those specific FIFOs. Cells coming from SAR (ATMizer) chip 526 are converted to a 16-bit format before they are written into FIFO 518. During this process, SOC alignment is also checked. If the SOC alignment is not right, the cell will be discarded and a SOC search will start until proper SOC alignment is found from the PHY chips.

Egress module 304 further periodically polls the ATZ and the Cubit chip FIFOs to see if they have a cell available to send to the ALM 502. This is done in a round-robin fashion. If they have a cell to send, the egress module 302 reads the first three words of the cell from the corresponding FIFO. Then it will start to respond to the UTOPIA addresses coming from the ALM as cells are available (only to the address zero). Cell transfer starts as soon as ALM 502 enables the cell read. At the same time, the rest of the cell from the FIFO is clocked in.

As shown in FIG. 5, egress module 304 includes a data path module 504, an ALM RX Utopia controller module 514, an ATZ FIFO controller

-18-

module 506, Cubit FIFO controller modules 508, 510, and a reset and CPU interface module 512.

Data path module 504 receives cells from the FIFOs 518, 520, 522 and performs cell multiplexing for the ALM 502. Selection signals come from the ALM RX Utopia controller module 514.

The ALM RX Utopia controller module 514 interfaces with the ALM's egress UTOPIA bus control signals portion. It also receives cell space available signals from the FIFO control modules 506, 508, 510. During normal operation, ALM 502 continuously polls the addresses of the single PHY device on the Utopia address bus. ALM RX Utopia controller 514 also continuously polls the FIFO cell availability signals coming from FIFO controller modules 506, 508, 510. If they indicate that there is a cell available, the corresponding PHY chip will be scheduled for the next transfer. As soon as the current cell transfer operation is completed, the first three words of the next cell will be read from the FIFO. After this, ALM RX Utopia controller 514 starts to respond to the ALM's polling by issuing cell space available signals only for the device which was previously scheduled.

If ALM 502 asserts the read enable signal, ALM RX Utopia controller module 514 will start to send the cell to the Utopia bus. At the same time it will start to read the rest of the cell which was waiting in the FIFO. During this process, the ALM RX Utopia controller module 514 assumes that cells were written into the FIFO with proper SOC alignment (this is the responsibility of FIFO controller module 506, 508, 510). But if something goes wrong and a cell comes out of the FIFO with a wrong alignment, this module will

-19-

generate an SOC Error signal to the Reset module. The ALM RX Utopia controller module 514 responds to single UTOPIA addresses of zero.

ATZ FIFO controller module 506 runs at the ALM's 50 MHz clock 550 and keeps the cell counter for the ATZ FIFO and provides interfaces to the ALM Utopia bus. If SAR (ATMizer) chip 526 has a cell to send and the corresponding FIFO cell counter indicates that there is room in FIFO 518, this module starts to read the cell from SAR (ATMizer) chip 526. During this process, it checks the SOC alignment. If SOC is not aligned with the cell, it will discard the cell and start searching for properly SOC aligned cells. If SOC is aligned properly, this module will do the byte to word conversion by using external staging registers and will write the cells in to the FIFO 518. It will also generate a cell available signal to the ALM RX Utopia controller module 514.

Cubit A and B FIFO controller modules 508, 510 respectively manage Cubit pro 524-A and 524-B chips running at 40 MHz clock 552. They keep the cell counters for the FIFOs and provide interfaces to the Cubit 524-A and Cubit 524-B 16-bit cell outlet buses.

If one of the Cubit chips 524-A and 524-B has a cell available and the FIFO cell counter indicates that there is room available in the respective FIFO 520, 522, these modules will start to read the cell from the corresponding Cubit chip. During this process they will check SOC alignment. If SOC is not aligned with the cell, they will discard the cell and start searching for properly SOC aligned cells. If SOC is aligned properly, they will write the cells into the corresponding FIFO. They will also

-20-

generate a cell available signal to the ALM RX Utopia controller module 514.

Reset and CPU interface module 512 generates a reset sequence for the following conditions:

(1) Missing SOC signal coming from ATMizer or Cubit FIFOs for cells going to data path module 504. This is an error condition on FIFO control, in which case the unaligned cell will be discarded.

(2) ATZ or Cubit chip FIFO overrun condition detected.

Normally, the FIFOs can not overflow. When they are almost full, PHY chips will not be allowed to send more cells to them using the Utopia control signals.

(3) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

During the reset sequence the Utopia interfaces are blocked by not allowing any cell operation. Internal logic is reset and the external FIFOs are flushed by generating proper reset signals for them.

Reset and CPU interface module 512 also provides the following pins for access by an external CPU (via CPU interface 516):

(1) "Main_reset_" is an asynchronous active low input signal. It has to be asserted for a minimum period of 1 ms. This will initiate a reset sequence, which resets the chip, indirectly resets the FIFOs and stops the Utopia interfaces. After main reset is de-asserted, it will take 32 clocks to go to operational mode. At the end of the sequence, pin "Err_stat" will be set as an acknowledgment, which signal needs to be cleared using the pin "Clear_err_cond".

-21-

(2) "Err_stat" is an active high output signal. It will be latched for the following conditions.

(a) Missing SOC signal for cells coming from the ATMizer or Cubit chip FIFOs to the egress module. This is an error condition on FIFO control, not a missing SOC on the ALM Utopia side to FIFO, in which case the unaligned cell will be discarded and hardware will realign itself to the in-coming cells by searching for cells with a proper SOC indicator.

(b) FIFO overrun condition detected. Normally the FIFOs should not overflow, and so this can only be a controller error. When the FIFOs are almost full, PHY chips will not be allowed to send more cells to them.

(c) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

When these conditions (a), (b) and (c) are detected, an auto initialization reset sequence is initiated. This will reset the chip and the FIFOs. This will also hold the Utopia interfaces in the blocking state, in which they will not be allowed to perform any cell transfers until the reset sequence is completed.

This signal can be cleared with the signal "clear_err cond".

(3) "Clear_err_cond" is an asynchronous input signal. It will clear the "Err_stat" output pin.

(4) "FIFO_Threshold" is an asynchronous input signal. It will determine the depth of the FIFOs for the Cubit chips. This value has to be set up initially and should not be changed during normal operation.

0 = Threshold is 8 cells

-22-

1 = Threshold is 144 cells

After power up, this module comes up in the reset mode (since Main reset is an active low signal) and stays in the reset mode until Main reset is deasserted. Before this signal is de-asserted, the ALM, Cubit and ATMizer chips need to be initialized to avoid any initial cell loss.

FIG. 6 further illustrates a narrowband card 202, 204 that can be included in the switch apparatus 102 illustrated in FIG. 2. It includes a cell bus controller 602 that communicates with two Cubit chips 628-A and 628-B, and a virtual circuit (VC) controller 608. The VC controller 608 further communicates with a packet controller 618, an ATM cell controller 620, and a voice/fax controller 626 via a shared bus 616. The cell bus controller 602 further communicates with first AAL1 SAR chip 604 and second AAL1 SAR chip 606.

The cell bus controller 602 is preferably implemented as a FPGA and provides five Utopia interfaces - - between the two Cubit chips 628-A and 628. B, the first AAL1 SAR chip 604, the second AAL1 SAR chip 606 and the VC controller 608. The cell bus controller 602 plays the role of an ATM layer multiplexer device providing interfaces to the shared buses 306-A and 306-B of the local switch module 206 from multiple PHY devices with different priorities. That is, the cell bus controller 602 multiplexes ATM cells from the low-speed interfaces connected to the NMC with the high-speed port of the ATM cell switching fabric, which the NMC shares with other NMCs via the shared bus of the local switch module 206.

The first and second AAL1 SAR chips 604, 606 are, for example, PMC73121 AALigator II chips from PMC-Sierra and are programmed to be in the single PHY mode. Between the SAR chips and the cell bus controller

-23-

there are FIFOs (not shown) which isolate the clocking domains. The SAR's Utopia runs at 33MHz. The FIFOs are, for example, SuperSync device IDT72261 from IDT of Santa Clara, CA. The first and second AAL1 SAR chips 604, 606 provide interfaces to ATM networks using T1/E1, T3/E3, and/or supported services.

VC controller 608 is preferably implemented by, for example, an L64364 ATMizer II+ from LSI. The VC controller's Utopia port is also configured to be in the single PHY mode. The Utopia clock runs at 40 MHz which is synchronous to the Cubit's Utopia clock. Also, the ATMizer is preferably configured to ignore parity on the Utopia bus.

As shown, VC controller 608 implemented by, for example an ATMizer chip, includes a SAR AALO/AAL5 engine 610, a plurality of VC queues 612, and a multi-service engine 614.

Generally, ATM cells received from, or to be sent to, the narrowband interfaces via shared bus 616 are stored in the VC queues 612. The rates at which the VC queues 612 are respectively serviced are determined by the quality of service for the particular flows to which the ATM cells stored in the VC queues belong. SAR engine 610 performs segmentation and reassembly functions for AAL5 frames and cell forwarding for AALO frames. It also runs a schedule-based service algorithm to determine which VC queue should be serviced (i.e., for transmitting a cell) at each cell time. VC queues 612 are implemented by a linked list of buffers. Multi-service engine 614 is a software module that performs packet translation if necessary for adapting flows for ATM cell switching in the switching fabric (e.g. FR to ATM network or service interworking), service functions based on header lookup, flow to VC mapping, and queuing of cells to the

-24-

appropriate per-VC queue 612. An example of a method and apparatus that can be used to implement the management of quality of service requirements for various flows according to the invention is described in co-pending U.S. Provisional Appln. No. 60/_____ Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference.

Shared bus 616 is preferably a PCI bus adapted for transfers of 1 Gbps.

Packet controller 618 interfaces between packets and frames that are input/output via attached IP/Frame Relay networks and ATM cells that are input/output via cell bus controller 602 and converted to packets and frames by VC controller 608. Preferably, it includes a HDLC controller (such as a PEB20324 from Siemens, for example) which performs HDLC functions such as bit stuffing/ unstuffing, CRC checking, etc. Packets and frames received by packet controller 618 from attached IP / Frame Relay networks are processed by the HDLC controller and sent to multi-service engine 614 via shared bus 616 for conversion to ATM cells and queuing in per-VC queues 612. Conversely, ATM cells destined for attached IP / Frame Relay networks are converted to packets by multi-service engine 614 and forwarded to packet controller 618 via shared bus 616, which then immediately forwards them to the appropriate IP / Frame Relay network port.

ATM cell controller 620 forwards ATM cells that are input from attached ATM networks to multi-service engine 614 via shared bus 616 and forwards ATM cells destined to attached ATM networks that are received

-25-

from the switch fabric via cell bus controller 602 to the attached ATM networks.

As shown, ATM cell controller 620 preferably includes an ATM service engine 622 and an ATM PLCP 624. ATM service engine 622 performs dual leaky bucket UPC based on ATM Forum 4.0 and uses a VPI/VCI index into a table to find the corresponding VC queue when forwarding cells between multiservice engine 614 and attached ATM networks. ATM PLCP 624 performs ATM physical layer functions according to ITU-.432 and direct cell mapping to DS1 or E1 transmission systems according to ITU-T G.804.

Voice/fax controller 626 converts voice/fax data that is received via attached PSTN networks into packets that are then forwarded to multi-service engine 614 via shared bus 616, and likewise converts packetized voice/data that is destined for attached PSTN networks and is received from multi-service engine 614 via shared bus 616 into voice/fax data for forwarding over the attached PSTN networks. An example of a method and apparatus that can be used to implement the management of voice/fax flows according to the invention is described in co-pending U.S. Provisional Appln. No. 60/_____, Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference.

FIG. 7 further illustrates a voice/fax controller 626 that can be included in the narrowband line card 202, 204 illustrated in FIG. 6. As shown, it preferably includes a digital signal processor (DSP) service engine 702, a DSP controller 706, and a set of DSPs 708-1...708-N. DSP

-26-

service engine 702 is preferably implemented by a LSI L64364 ATMizer II+. DSPs 708 are preferably implemented by TMS320C6201 chips.

As shown, DSP controller 706 provides the interface between the DSP service engine and the DSPs. DSP service engine 702 communicates with DSP controller 706 via a UTOPIA data path interface from which it receives, decodes and executes messages. DSP service engine 702 also controls the DSP controller 706 via its secondary port interface (described below).

DSP service engine 702 operates in slave mode, with cell data outputs from its transmit FIFO (not shown) going to the UTOPIA master, and cell data inputs from DSP controller 706 master being input to its receive FIFO 704. DSP controller 706 accesses the DSPs via its host port interface (described below). DSP service engine is mainly responsible for performing media type adaptation of voice/fax flows received or sent over PSTN networks connected to switch apparatus 102.

A DSP service engine 702 that can be adapted for use in the present invention is described in co-pending U.S. Provisional Application No.

_____, Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference. Briefly, however, TDM-based voice/fax data streams received by voice/fax controller 626 from attached PSTN networks (via a time slot interchanger known in the art, for example) are packetized by the DSPs and stored in DSP output queues (not shown). DSP controller 706 forwards the voice/fax packets from the DSP output queues to DSP service engine 702. DSP service engine 702 then assembles the packets in accordance with the egress media type for the voice/fax connection. If the

-27-

egress media type is IP, for example, the DSP service engine adds a RTP, UDP, and IP header to the packets and forwards them to multi-service engine 614 via shared bus 616. Conversely, packetized voice/fax data streams are received from multi-service engine 614 via shared bus 616. DSP service engine 702 then converts the packets of the ingress media type into voice/ fax data. For example, if the ingress media type was IP, DSP service engine 702 reads the RTP, UDP and IP headers of the packets and writes the packets to the appropriate DSP input queues (not shown). The DSPs convert the packetized data into voice/fax streams for output via attached PSTN networks.

FIG. 8 further illustrates a DSP controller 706 such as that illustrated in FIG. 7. As shown, it includes a cell receive block 802, a cell transmit block 804, a command processor 808, a secondary port interface (SPI) block 806, a host port interface (HPI) block 814, and block read unit 812, and a block write unit 810.

The cell receive block 802 interfaces with the DSP service engine 702 receive port. DSP controller 706 acts as Utopia master and DSP service engine 702 as slave. The cell receive block 802 is responsible for loading the messages into internal RAM buffers (not shown), which provide storage for two cells. Command processor 808 issues a load command to cell receive block 802, and cell receive block 802 generates a ready signal when the corresponding buffer contains valid data. The command processor 808 then decodes and executes the contents of the current buffer while cell receive block 802 is loading the next buffer.

Cell receive block 802 checks the parity of the receive data bus and compares it to the incoming parity bit. The parity error bit in Interrupt 0

-28-

vector register is set and the DSP controller 706 asserts an interrupt signal when a parity error is detected.

Cell receive block 802 saves the parity bits for the header word, which bits are looped back for a Block Read operation when the header word is written to the external FIFO.

The cell buffers (not shown) are implemented using embedded 256 by 8 bit wide RAMS. The command processor 808 is able to overlap the reading of the buffers with transferring of data to the host port interface 814 by using a "prefetch" mechanism. The command processor 808 asserts a next word load signal and cell receive block 802 automatically reads the next 4 bytes and form a 32 bit word at the buffer data outputs.

Cell transmit block 804 interfaces to DSP service engine 702's transmit port. The block read unit 812 directly writes cell data to external FIFO 704 while the almost full flag remains de-asserted.

Immediately after reset, cell transmit block 804 waits until the block read unit 812 has programmed the external FIFO almost full and almost empty flags. When the external FIFO 704 contains at least one full cell, the almost empty flag will be de-asserted. When the DSP service engine 702 asserts a signal indicating that a cell is available the cell transmit block 804 will proceed to generate the control signals to read the data out of the external FIFO 704 and into the DSP service engine 702. The cell transmit block 804 maintains an octet counter to keep track of the words read out of the external FIFO and to generate a start of cell signal.

A read enable signal is input to the external FIFO 704, which signal is delayed one clock, thus indicating valid data on the input pins. The external FIFO almost empty flag is valid on the second clock

-29-

following a read cycle and it is registered at the input of the DSP controller 706. The cell transmit state machine 804 samples the internal almost empty flag three clock cycles after the last byte of a cell have been read out of the external FIFO to accommodate the latencies.

SPI block 806 communicates with the secondary port interface of the DSP service engine 702 so as to provide direct access to DSP memory locations, HPI block registers, DSP reset control, reading and clearing of the Interrupt 0 vector register, and reading of the Interrupt 1 vector register.

The SPI runs asynchronously to the DSP controller 706. SPI block 806 detects the falling edge of an access signal from the SPI to initiate a secondary port access to the DSP controller 702, and asserts a ready signal when the cycle is completed. Two address lines and a write enable signal are used to provide access to four addressable locations within the DSP controller 702: a read/write control register, a read/write data register, the Interrupt 0 vector register, and the Interrupt 1 vector register.

The read/write control register is loaded first and is used to set up the starting addresses for reads and writes, as well as selecting one of the DSPs. The OpCode field of the read/write control register and the write enable signal determine the behavior of the read/write operation. An OpCode value of "100" causes a read or write operation to the selected address of the selected DSP to be performed. An OpCode value of "101" causes a write operation into the DSP Reset register. OpCode values of "000" through "011" cause the indicated HPI register of the selected DSP to be accessed for diagnostic purposes.

-30-

Selection of the Interrupt 0 and 1 vector registers allows the contents of these registers to be viewed.

Command processor 808 issues buffer load requests to cell receive block 802. When the cell receive block 802 has loaded 64 bytes into a cell buffer it will assert the corresponding buffer ready flag. The command processor 808 decodes the message type field of the header word and asserts a start signal to the block write unit 810 or the block read unit 812. The command processor 808 issues a buffer load request to the alternate buffer while the block write or block read units 810, 812 are operating on the data from the first cell buffer. The command processor 808 also muxes the data and control signals between the cell

When block write unit 810 receives a start signal from the command processor 808, it loads the cell data message length field into a message counter register, loads the DSP ID field into a register, and asserts a signal to "prefetch" the read/write DSP address word from the buffer. It asserts an initial request signal to the HPI to let it know that this transfer will load the address register as well as perform one 32 bit data write cycle to the corresponding DSP.

Block write unit 810 waits for an acknowledge signal to be asserted by the HPI. When the HPI acknowledges that the DSP's HPIA register has been loaded, the block write unit 810 will "prefetch" the first data word to be written and again wait for the acknowledge signal. For subsequent write cycles the block write unit 810 asserts subsequent request signals.

Block write unit 810 uses the message counter to keep track of how many 4 byte words are to be written, and it also uses a word counter

-31-

to keep track of the number of 4 byte words remaining in this cell. It then determines if another cell is needed to complete the block write operation. When all the words in the current cell have been written, a signal is asserted indicating this. It should be noted that the first cell has a payload of 14 words (56 bytes) and subsequent cells have a payload of 15 words (60 bytes).

When block write unit 810 is finished writing all the words indicated by the message counter, it generates a signal to the SPI to assert the End of Block interrupt, which signal further identifies which bit to assert in the Interrupt 0 vector register.

Block read unit 812 drives the data and control signals to the TX FIFO 704. After reset it generates the necessary control signals to program the almost empty flag to be asserted when less than one cell is in the TX FIFO (64 bytes) and the almost full flag is asserted when less than one cell (64 bytes) can be written. When block read unit 812 receives a start signal from the command processor 808, it loads the cell data message length field into a message counter register, loads the DSP ID field into a register, and stores the entire cell data into the header register.

The block read unit then checks the TX FIFO almost full flag. If there is room for an entire cell it writes the header register and asserts a signal to "prefetch" the read/write DSP address word from the buffer. It asserts the initial request signal to the HPI to let it know that this transfer will load the address register as well as perform one 32 bit data read cycle to the corresponding DSP. If there is no room in TX FIFO 704 for an entire cell, the block read unit 812 "aborts" and the cell is not read.

-32-

Block write unit 810 waits for an acknowledge signal to be asserted by HPI 814. When the HPI acknowledges that the DSP HPIA register has been loaded, and the first 4 byte word has been read, block read unit 812 writes the data one byte at a time to the TX FIFO 704. Block read unit 812 uses the word counter to determine when a new cell is started and when to write the header register into the TX FIFO.

When the number of words to read does not fill an integral number of cells, block read unit 812 starts the FILL state machine. Its function is to take over and "fill" up the cell with data while the block read unit 812 indicates to the command processor 808 that it is done.

Host port interface 814 is responsible for arbitrating between the secondary port requests and the block read or block write unit requests. Priority is given to SP requests; however, when the block read or block write units 810, 812 assert their initial request signals, the address and data portions of the transfers are not interruptible by the SPI 806.

The arbiter will first check for SPI requests. If an SP initial request signal is asserted, the opcode bits are checked to determine if the secondary port will perform a memory read/ write ("100") or if direct access to the DSP HPI registers is required. For direct access the opcode bits will drive the control bits. For memory transfers, the state machine will indicate HPIA during the address cycle and HPID during data cycle. If the secondary port performs a memory access, the DSP ID is checked against the DSP ID of the last block read or block write transfer and the HIT flag is set if the secondary port access is to the same DSP.

-33-

If no SP request is asserted, then the block write and block read request signals are checked. If an initial request is made or the HIT flag is true, then the HPIA address register is loaded and a read or write cycle is performed and the HIT flag is cleared. If a data request transfer is made and the HIT flag is false then there is no need to reload the HPIA register.

In the above description of an example of a narrowband card 202, 204 according to the present invention, egress channel adaptation is simply the reverse process of ingress channel adaptation. That is, if necessary, ATM cells forwarded to the narrowband card via another narrowband card (via the shared bus of local switch module 206) or from a broadband card (via switch fabric 208 and local switch module 206), are converted into packets or frames as appropriate for forwarding on an attached non-ATM network.

It should be noted that broadband card BSC1 210 also preferably includes similar functionality for adapting ingress and egress broadband flows for switching, for example by enqueueing ATM cells communicated or to be communicated over cell switching fabric 208 in virtual circuit queues, and transforming, if necessary, between such ATM cells and any other media type in accordance with the network on the broadband network connections of BSC1. For example, if an egress media type is Ethernet, the BSC1 reassembles the ATM cells received from cell switching fabric 208 into an AAL5 frame, performs service functions, and enqueues the frame to the corresponding Ethernet MAC Tx queue. Conversely, if an ingress media type is Ethernet, the BSC 1 segments

-34-

AAL5 frames from the corresponding Ethernet MAC Rx queue into ATM cells for switching via cell switching fabric 208.

It should be further apparent that the present invention allows for switching flows of data between two broadband networks or two narrowband networks, as well as allowing for switching flows between a broadband network and a narrowband network as is more explicitly described above.

Although the present invention has been described in detail with reference to the preferred embodiments thereof, those skilled in the art will appreciate that various substitutions and modifications can be made to the examples described herein while remaining within the spirit and scope of the invention as defined in the appended claims.

-35-

What is claimed is:

1. A carrier class switch apparatus integrated in a single switching platform comprising:
 - a switching fabric adapted to switch traffic between a plurality of broadband switching ports;
 - a broadband interface coupled to one of the plurality of broadband switching ports; and
 - a local switch module coupled to another one of the plurality of broadband switching ports and to one or more narrowband interfaces.
2. An apparatus according to claim 1, wherein the broadband interface is adapted to be further coupled to one or more broadband networks.
3. An apparatus according to claim 1, wherein the one or more narrowband interfaces are each adapted to be further coupled to one or more narrowband networks.
4. An apparatus according to claim 1, further comprising a switch control card coupled to the switching fabric, the broadband interface and the local switch module, wherein the switch control card is adapted to route and manage virtual circuit connections between the plurality of broadband switching ports in accordance with a quality of service requirement.

-36-

5. An apparatus according to claim 1, wherein the local switch module is adapted to cause the one or more narrowband interfaces to emulate the 5 broadband interface.

6. An apparatus according to claim 1, wherein the local switch module comprises:

a shared bus coupled to the one or more narrowband interfaces;

an ingress module coupled to the shared bus and the switching fabric that is adapted to forward traffic from the switching fabric to appropriate narrowband interface; and

an egress module coupled to the shared bus and the switching fabric that is adapted to forward traffic from the narrowband interfaces to the shared bus.

7. An apparatus according to claim 6, wherein the shared bus is comprised of a plurality of buses configured for load sharing mode.

8. An apparatus according to claim 6, wherein the shared bus is comprised of a plurality of buses configured for redundancy mode.

9. An apparatus according to claim 6, wherein the ingress module includes:

-37-

one or more data paths between the switching fabric and the shared bus, the data paths adapted to buffer traffic between the switching fabric and corresponding narrowband interfaces; and

a controller module adapted to communicate with the switching fabric and to map traffic received from the switching fabric to the one or more data paths.

10. An apparatus according to claim 6, wherein the egress module includes:

one or more data paths between the switching fabric and the shared bus, the data paths adapted to buffer traffic between the switching fabric and corresponding narrowband interfaces; and

a controller module adapted to communicate with the switching fabric and to map traffic received from the narrowband interfaces to the switching fabric.

11. An apparatus according to claim 9, wherein the egress module includes:

one or more data paths between the switching fabric and the shared bus, the data paths adapted to buffer traffic between the switching fabric and corresponding narrowband interfaces; and a controller module adapted to communicate with the switching fabric and to map traffic received from the narrowband interfaces to the switching fabric.

12. An apparatus according to claim 6, wherein at least one of the narrowband interfaces comprises:

-38-

at least one network interface adapted to communicate traffic with a narrowband network;

a virtual circuit queue for buffering traffic associated with the narrowband network; and

a cell bus controller coupled to the virtual circuit queue and the shared bus, the cell bus controller adapted to multiplex the narrowband network traffic buffered in the virtual circuit queue with the broadband switching port of the switching fabric to which the shared bus is coupled.

13. An apparatus according to claim 11, wherein at least one of the narrowband interfaces comprises:

at least one network interface adapted to communicate traffic with a narrowband network;

a virtual circuit queue for buffering traffic associated with the narrowband network; and

a cell bus controller coupled to the virtual circuit queue and the shared bus, the cell bus controller adapted to multiplex the narrowband network traffic buffered in the virtual circuit queue with the broadband switching port of the switching fabric to which the shared bus is coupled.

14. An apparatus according to claim 12, wherein the at least one narrowband interface further includes a multi-service engine that translates between a first media type associated with the narrowband network traffic and a second media type of the switching fabric.

-39-

15. An apparatus according to claim 14, wherein the first media type is IP and the second media type is ATM.

16. An apparatus according to claim 14, wherein the first media type is ATM and the second media type is ATM.

17. An apparatus according to claim 14, wherein the first media type is frame relay and the second media type is ATM.

18. An apparatus according to claim 14, wherein the first media type is voice and the second media type is ATM.

19. An apparatus according to claim 14, wherein the first media type is fax and the second media type is ATM.

20. An apparatus according to claim 2, wherein the broadband interface includes a multi-service engine that translates between a first media type associated with broadband network traffic and a second media type of the switching fabric.

21. An apparatus according to claim 14, wherein the broadband interface is adapted to be further coupled to one or more broadband networks, and the broadband interface includes a second multi-service engine that translates between a third media type associated with broadband network traffic and the second media type of the switching fabric.

-40-

22. A carrier-class switch apparatus integrated in a single switch platform comprising:

a switching fabric adapted to switch ATM cells between a plurality of broadband ports;

a broadband interface adapted to be coupled between one of the broadband ports and a broadband network, the broadband interface including a broadband multi-service engine that is adapted to translate between ATM cells switched via the switching fabric and an egress media type associated with, broadband network traffic;

a local switch module coupled to another of the broadband ports, the local switch module including a shared bus; and

a narrowband interface adapted to be coupled between the shared bus of the local switch module and a narrowband network, the narrowband interface including a narrowband multi-service engine that is adapted to translate between ATM cells switched via the switching fabric and an ingress media type associated with narrowband network traffic.

23. An apparatus according to claim 22, wherein the ingress media type is IP.

24. An apparatus according to claim 22, wherein the ingress media type is ATM.

25. An apparatus according to claim 22, wherein the ingress media type is frame relay.

-41-

26. An apparatus according to claim 22, wherein the ingress media type is voice.

27. An apparatus according to claim 22, wherein the egress media type is IP.

28. An apparatus according to claim 22, wherein the egress media type is ATM.

29. An apparatus according to claim 22, wherein the egress media type is frame relay.

30. An apparatus according to claim 22, wherein the egress media type is voice.

31. An apparatus according to claim 26, wherein the egress media type is IP.

32. An apparatus according to claim 26, wherein the egress media type is ATM.

33. An apparatus according to claim 26, wherein the egress media type is frame relay.

-42-

34. An apparatus according to claim 26, wherein the egress media type is voice.

35. A method for switching traffic between networks in an integrated carrier-class switching platform, the method comprising:

- (a) communicating broadband network traffic via a broadband interface;
- (b) adapting the broadband network traffic for switching via a broadband switching port of a switching fabric;
- (c) communicating narrowband network traffic via at least one narrowband interface;
- (d) multiplexing the narrowband network traffic on a shared bus; and
- (e) adapting the narrowband network traffic for switching via another broadband switching port of the switching fabric so that the at least one narrowband interface emulates the broadband interface.

36. A method according to claim 35, wherein the steps of adapting the broadband and narrowband network traffic for switching each include converting the network traffic from a first media type to a second media type:

37. A method according to claim 36, wherein the second media type is ATM.

-43-

38. A method according to claim 37, wherein the first media type is IP.

39. A method according to claim 37, wherein the first media type is ATM.

40. A method according to claim 37, wherein the first media type is voice.

41. A method according to claim 37, wherein the first media type is frame relay.

42. An apparatus for switching traffic between networks in an integrated carrier-class switching platform, the apparatus comprising:
means for communicating broadband network traffic via a broadband interface;
means for adapting the broadband network traffic for switching via a broadband switching port of a switching fabric;
means for communicating narrowband network traffic via at least one narrowband interface;
means for multiplexing the narrowband network traffic on a shared bus; and
means for adapting the narrowband network traffic for switching via another broadband switching port of the switching fabric so that the at least one narrowband interface emulates the broadband interface.

-44-

43. An apparatus according to claim 42, wherein the means for adapting the broadband and narrowband network traffic for switching each include means for converting the network traffic from a first media type to a second media type.




FIG. 1

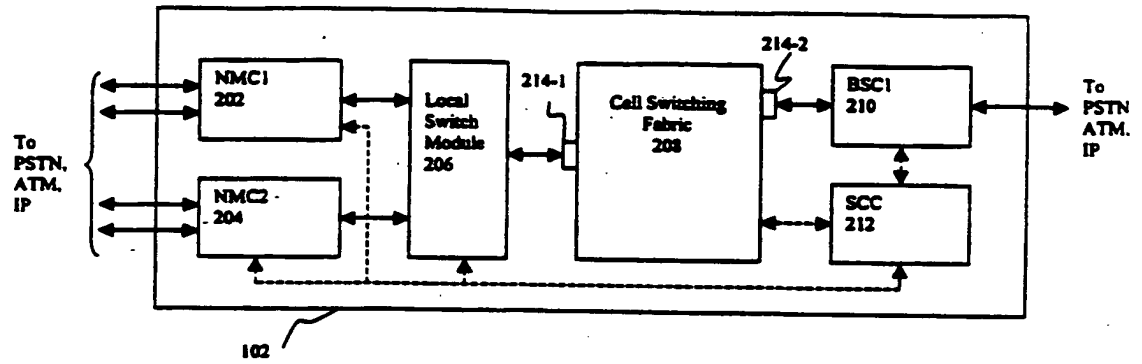


FIG. 2




FIG. 3

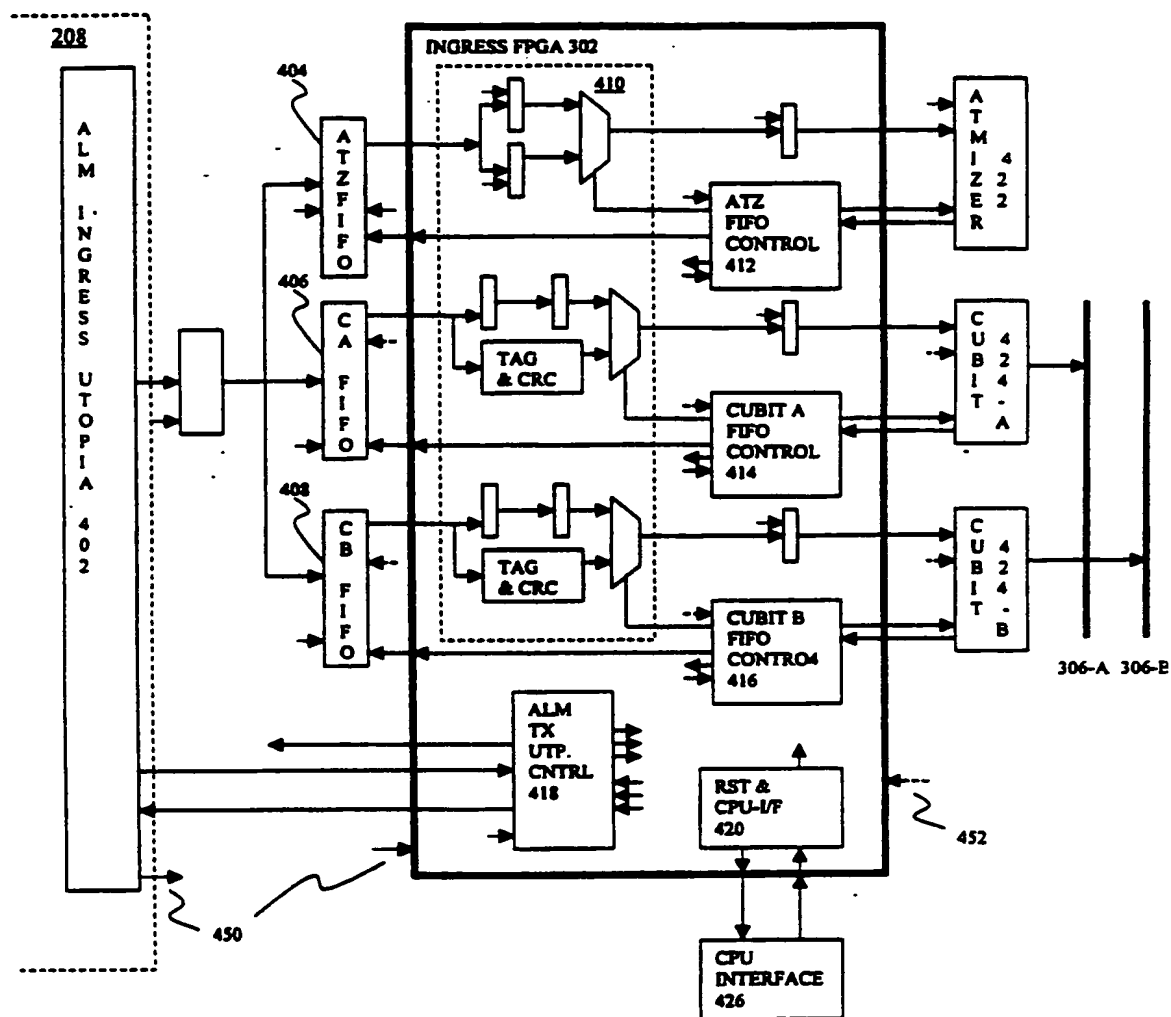


FIG. 4

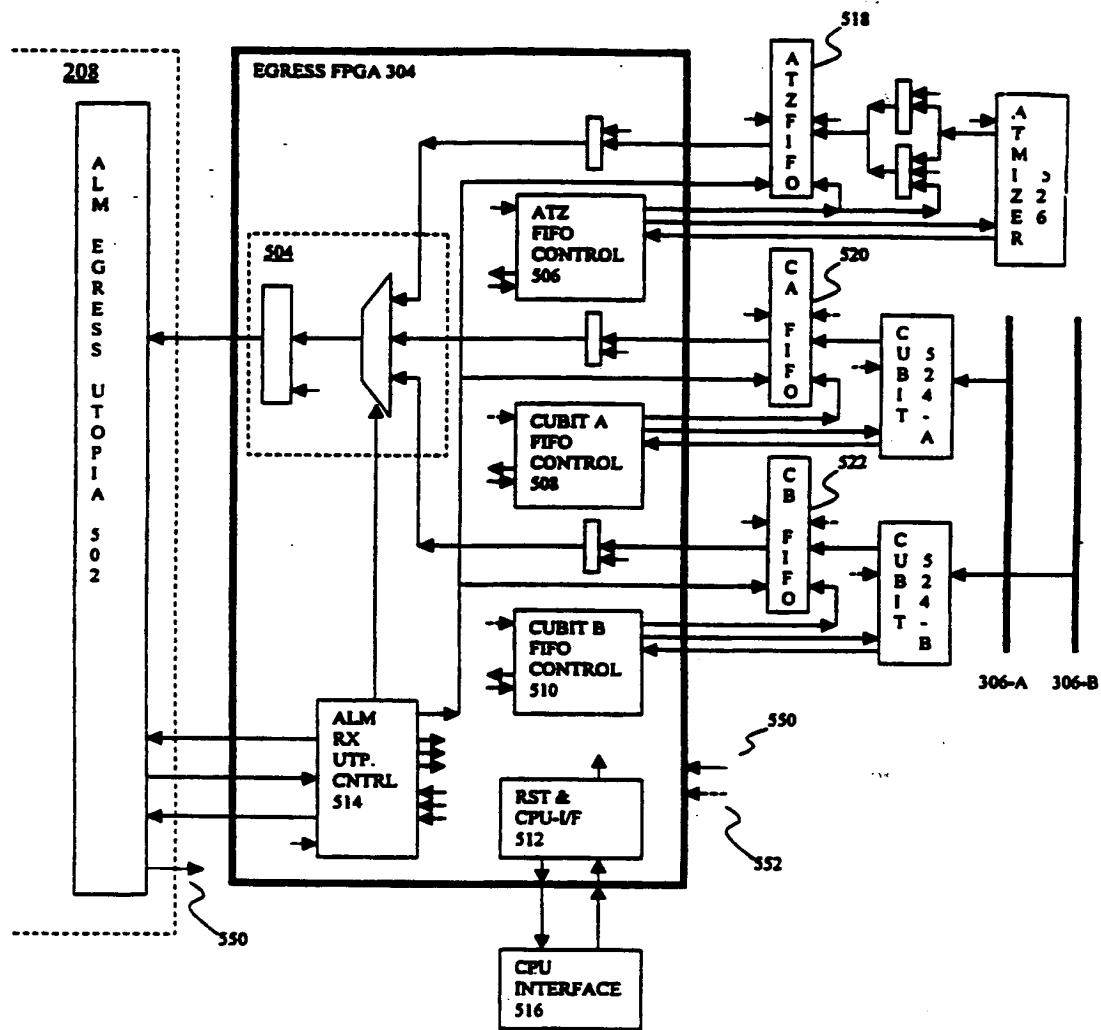


FIG. 5

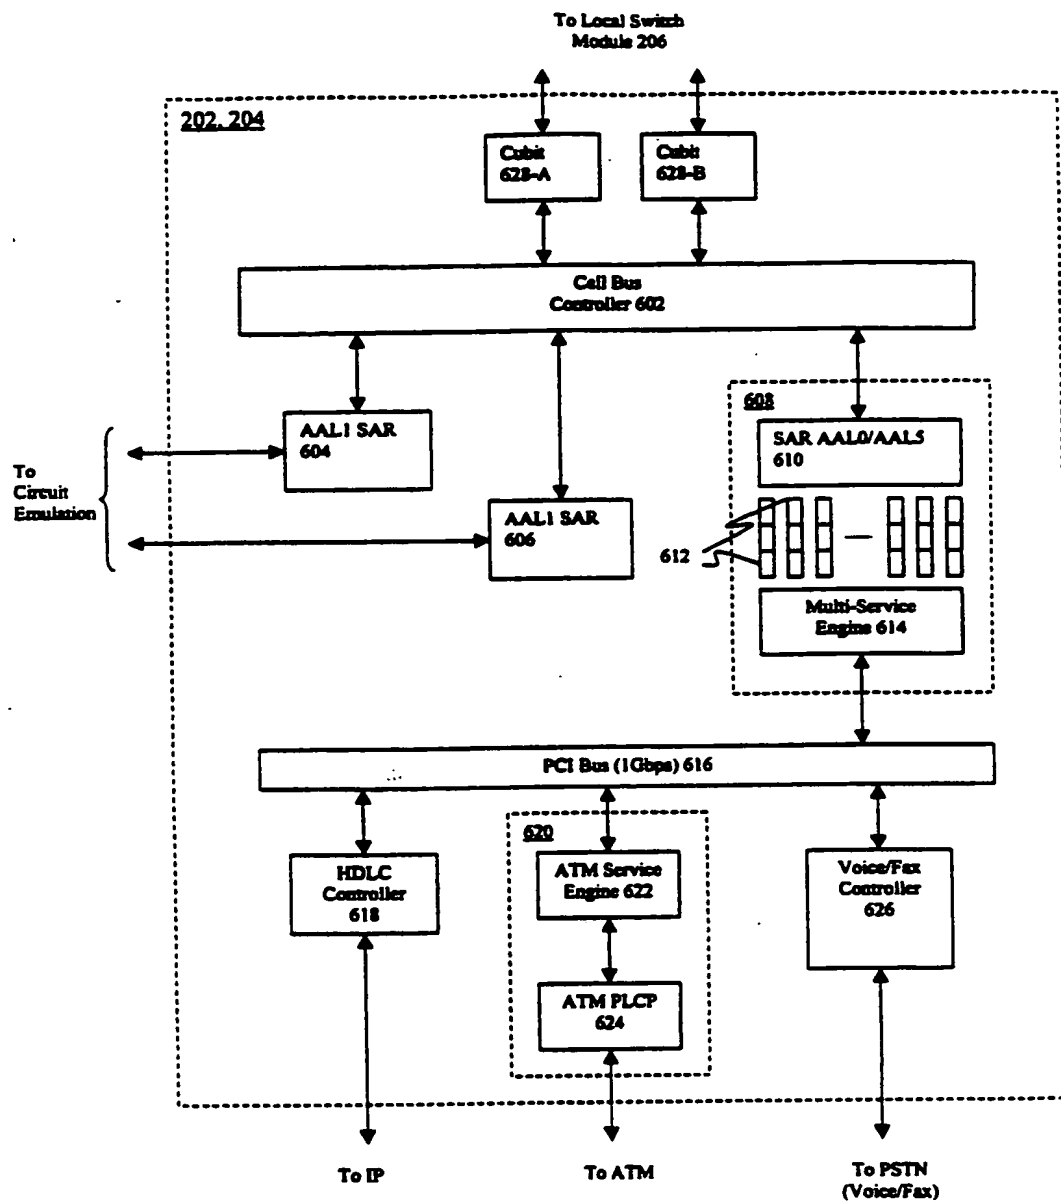


FIG. 6




FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/06481

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/64 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US 5 809 022 A (RUNYON JAMES PHILIP ET AL) 15 September 1998 (1998-09-15) figures 4,7	1,22,35, 42,43 2-21, 23-34, 36-41
P,X	WO 00 02410 A (DEML REINHARD ;SIEMENS AG (DE); WAHLER JOSEF (DE)) 13 January 2000 (2000-01-13) figure 2	1,22,35, 42,43
X A	WO 94 03004 A (ITALTEL SPA ;CANATO LUIGI (IT); GALLASSI GIORGIO (IT); MORGANTI MI) 3 February 1994 (1994-02-03) figure 2	1 2-43
	--- -/-	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

26 June 2000

Date of mailing of the international search report

04/07/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5816 Patentaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Meurisse, W

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/06481

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>PERES M: "SCHALTEN WIE IN DER FORMEL 1" ELEKTRONIK, DE, FRANZIS VERLAG GMBH. MUNCHEN, vol. 47, no. 14, July 1998 (1998-07), pages 68-78, XP000669736 ISSN: 0013-5658 page 74, left-hand column, paragraph 4 -right-hand column, paragraph 4 page 78, left-hand column, paragraph 4 -right-hand column, paragraph 3 -----</p>	1-43
A	<p>US 5 781 320 A (BYERS CHARLES CALVIN) 14 July 1998 (1998-07-14) figure 3 -----</p>	6, 22, 35, 42

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/US 00/06481

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5809022 A	15-09-1998	EP 0797373 A JP 10065742 A	24-09-1997 06-03-1998
WO 0002410 A	13-01-2000	DE 19829822 A	05-01-2000
WO 9403004 A	03-02-1994	IT 1259036 B DE 69309471 D DE 69309471 T EP 0653132 A	11-03-1996 07-05-1997 06-11-1997 17-05-1995
US 5781320 A	14-07-1998	NONE	

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
14 September 2000 (14.09.2000)

PCT

(10) International Publication Number
WO 00/54469 A1

(51) International Patent Classification⁷: **H04L 12/64**,
12/56

(21) International Application Number: **PCT/US00/06481**

(22) International Filing Date: **10 March 2000 (10.03.2000)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:
60/123,798 11 March 1999 (11.03.1999) US
09/405,993 27 September 1999 (27.09.1999) US

(71) Applicant: **CISCO TECHNOLOGY, INC.** [US/US];
170 West Tasman Drive, San Jose, CA 95134 (US).

(72) Inventors: **KAO, Ruey**; 13502 Myren Drive, Saratoga,
CA 95070 (US). **SUE, Cherng-Ren**; 6066 Prince Drive,
San Jose, CA 95129 (US). **YILDIRIM, Cem**; 6405 Samar
Drive, San Jose, CA 95119 (US).

(74) Agents: **VINCENT, Lester, J. et al.**; Blakely, Sokoloff,
Taylor & Zafman, 7th floor, 12400 Wilshire Boulevard, Los
Angeles, CA 90025-1026 (US).

(81) Designated States (*national*): AE, AL, AM, AT, AU, AZ,
BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK,
DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID,
IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT,
LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL,
PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ,
UA, UG, UZ, VN, YU, ZA, ZW.

(84) Designated States (*regional*): ARIPO patent (GH, GM,
KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent
(AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent
(AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU,
MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM,
GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:


— with international search report

(48) Date of publication of this corrected version:
25 October 2001

(15) Information about Correction:
see PCT Gazette No. 43/2001 of 25 October 2001, Section
II

[Continued on next page]

(54) Title: **MULTIMEDIA AND MULTIRATE SWITCHING METHOD AND APPARATUS**



(57) Abstract: A switch apparatus and method according to the invention implements a three stage switching process. Various types of media streams presented to the switch apparatus by the broadband and narrowband connections are adapted for switching by being converted to ATM cells and enqueued in corresponding virtual circuit (VC) queues. ATM cell switching is performed among the different cards based on the quality of service requirement for each virtual circuit. The switched ATM cells are then converted to the outgoing media types and outputted to the necessary broadband and narrowband connections. The switch apparatus and method is further adapted to perform rate shaping and traffic management so as to guarantee the quality of service for various media types (voice, video, data) and also minimize the traffic loss due to rate mismatch between narrowband and broadband connections during the burst period. By virtue of this implementation, the switch apparatus and method of the present invention can perform any-to-any media type switching.

WO 00/54469 A1

WO 00/54469 A1



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

-1-

MULTIMEDIA AND MULTIRATE SWITCHING METHOD AND APPARATUS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to carrier class switches, and in particular, to a method and apparatus for providing multimedia and multirate switching in an integrated carrier-class switching platform.

2. Description of the Related Art

Most conventional switches are adapted for switching traffic of a specific media such as voice switches (e.g. 4ESS, DMS200) and data switches (e.g. ATM/Frame Relay), or for switching traffic of a specific speed such as backbone switches (switching traffic among broadband interfaces, e.g. OC-12) and access switches (switching traffic among narrowband interfaces, e.g. DSO).

In particular, a backbone switch typically includes a switching fabric that switches between a certain number (e.g. 8 or 16) of high speed ports such as OC-12 ports. In contrast, low speed traffic is generally switched between low speed ports in an access switch having, for example, a shared bus architecture. For such low speed traffic to access one of the high-speed lines coupled to the switching fabric (and vice-versa), access switches generally include an uplink module to convert the low-speed traffic to high-speed traffic, and a separate OC-12 line is needed to transmit the converted low-speed traffic to the high-speed switching fabric. Such methods typically require one high speed port to be dedicated for each

-2-

low-speed uplink connection, thus wasting the available bandwidth on the high-speed port.

Moreover, traffic can typically be switched within one type of network. For example, ATM traffic can only be switched among ATM networks, IP traffic can only be switched among IP networks, etc.

Accordingly, there remains a need in the art for an integrated switching apparatus that provides for switching among both low speed ports (e.g. NxDSO) and high speed ports (e.g. OC-3). Moreover, such a switching apparatus should be able to minimize the traffic loss due to rate mismatch during the burst period.

There further remains a need in the art for a switch apparatus that provides for switching among different types of multimedia streams. Such a switching apparatus should also be able to guarantee the quality of service for the different media types, e.g. voice, video, data.

SUMMARY OF THE INVENTION

Accordingly, an object of the present invention is to provide a switching apparatus and methodology that permits traffic of any speed to be switched between ports of a single device.

Another object of the present invention is to provide a switching apparatus and methodology that permits traffic of any media to be switched in a single device.

Another object of the present invention is to provide a switching apparatus and methodology that permits traffic of any network to be switched in a single device.

-3-

Another object of the present invention is to provide a switching apparatus and methodology that minimizes traffic loss due to rate mismatch between narrowband and broadband connections during the burst period.

Another object of the present invention is to provide a switching apparatus and methodology that guarantees quality of service for flows of various media types.

To achieve these objects and others, the switch apparatus and method according to the invention implements a three stage switching process. Various types of media streams presented to the switch apparatus via broadband and narrowband flows (e.g. voice/fax call, video session, packet flow between source and destination ports, etc.) are converted to ATM cells and enqueued in corresponding virtual circuit (VC) queues. ATM cell switching is performed among the different cards based on the quality of service required for each virtual circuit. The switched ATM cells are then converted to the outgoing media types and outputted to the necessary broadband and narrowband flows. The switch apparatus and method are further adapted to perform rate shaping and traffic management so as to guarantee the quality of service for various media types (voice, video, data) and also minimize the traffic loss due to rate mismatch between narrowband and broadband connections during the burst period. By virtue of this implementation, the switch apparatus and method of the present invention can perform any-to-any media type switching.

-4-

BRIEF DESCRIPTION OF THE DRAWINGS

These and other objects and advantages of the present invention will become apparent to those skilled in the art after considering the following detailed specification, together with the accompanying drawings wherein:

FIG. 1 illustrates an implementation of a switch apparatus according to the present invention;

FIG. 2 is a block diagram of a switch apparatus according to the present invention;

FIG. 3 is a block diagram further illustrating an example of a local switch module that can be included in a switch apparatus of the present invention such as that shown in FIG. 2;

FIG. 4 is a block diagram further illustrating an example of an ingress module that can be included in a local switch module such as that shown in FIG. 3;

FIG. 5 is a block diagram further illustrating an example of an egress module that can be included in a local switch module such as that shown in FIG. 3;

FIG. 6 is a block diagram further illustrating an example of a narrowband service card that can be included in a switch apparatus of the present invention such as that shown in FIG. 2;

FIG. 7 is a block diagram further illustrating an example of a voice/fax controller module that can be included in a narrowband service card such as that shown in FIG. 6; and

FIG. 8 is a block diagram further illustrating an example of a DSP controller module that can be included in a voice/fax controller module such as that shown in FIG. 7.

-5-

DETAILED DESCRIPTION OF THE INVENTION

FIG. 1 illustrates an exemplary implementation of the present invention. As illustrated, a switch apparatus 102 according to the present invention provides, within a single device, the capability to switch flows of any media among LEC/CLECs, PSTNs, and ATM and IP networks.

An example of a switch apparatus 102 according to the present invention such as that shown in FIG. 1 is further illustrated in FIG. 2.

As shown in FIG. 2, switch apparatus 102 includes an ATM cell switching fabric 208 that switches ATM cell traffic between switch ports 214-1...214-N. Coupled to one of the switch ports is a broadband service card (BSC1) 210 for interfacing with a plurality of broadband connections. Coupled to a second one of the switch ports is a switch control card (SCC) 212. Coupled to another of the switch ports is a local switch module 206. Further coupled to the local switch module 206 is a plurality of narrowband line cards (NMC1, NMC2) 202, 204 for interfacing with a plurality of narrowband connections.

ATM cell switching fabric 208 is, for example, an ATLANTA chipset switch fabric having an 8x8 array of switch elements such as LUC4AS01 ATM Switch Elements made by Lucent Technologies of Allentown, PA. Such a switch fabric switches ATM traffic between eight switch ports 214-1...214-8 (only two such ports are shown in FIG. 2 for clarity). Switch ports 214 are preferably OC-12 or equivalent ports. An implementation of such an ATM cell switching fabric is described in Lucent Technologies Product Brief, "ATLANTA ATM Switch Core Chip Set," March 1997, the contents of which are incorporated herein by reference.

-6-

BSC1 210 provides an interface between one switch port 214-2 of the ATM cell switching fabric 208 and one or more broadband connections such as T3/E3, OC-3, and OC-12 lines and/or ports. Although only one broadband service card is shown, it should be apparent that there may be several.

SCC 212 contains functionality for establishing, routing, and managing virtual circuit connections between the ports of the switch apparatus. An example of an apparatus and method that can be used to implement such functionality is described in co-pending U.S. Provisional Appln. No. 60/_____, Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference.

Local switch module 206 provides an interface between one switch port 214-1 of the ATM cell switching fabric 208 and one or more narrowband line cards NMC1 202, NMC2 204. The switch port 214-1 of the ATM cell switching fabric 208 that is coupled to local switch module 206 is configured to control, for example, 16 multiPHY devices on the physical layer side. This can be implemented using, for example, a LUC4AU01 ATM Layer UNI Manager (ALM) from Lucent Technologies (not shown). Transfers of ATM cells between ATM cell switching fabric 208 (via ALM) and local switch module 206 are preferably performed via a 16-bit UTOPIA II interface (not shown). The local switch module 206 thus allows all the narrowband connections to share the bandwidth of one broadband connection. This improves the prior art solution of separately adapting one broadband connection for each narrowband interface.

-7-

NMC1 202 and NMC2 204 each provide an interface between local switching module 206 and one or more narrowband connections such as NxDSO, NxT1/E1, Ethernet, ISDN lines and/or ports. Although two narrowband interface cards are shown, it should be apparent that there may be one or several.

Generally, the switch apparatus illustrated in FIG. 2 implements a three stage switching process. Various types of media streams presented to the switch apparatus by the broadband and narrowband flows are adapted for switching between connected networks by being converted to ATM cells and enqueued in corresponding virtual circuit (VC) queues. ATM cell switching is performed among the different cards based on the quality of service required for each virtual circuit. The switched ATM cells are then converted to the outgoing media types and outputted to the necessary broadband and narrowband connections. The switch apparatus is adapted to perform rate shaping and traffic management so as to guarantee the quality of service for various media types (voice, video, data) and also minimize the traffic loss due to rate mismatch between narrowband and broadband connections during the burst period. By virtue of this implementation, the switch apparatus of the present invention can perform any-to-any media type switching as listed in Table 1.

-8-

Table 1 Switching Matrix

Input\Output	Voice/Fax/Video	Frame (FR, PPP)	ATM Cell	LAN
Voice/Fax/Video	Voice Switching	VoFR/VoIP	VoATM VoIP	VoIP
Frame (FR, PPP)	VoFR/VoIP	Frame Switching	FR/ATM Interworking	Encapsulation e.g. RFC 1490
ATM Cell	VoATM/VoIP	FR/ATM Interworking	Cell Switching	Encapsulation e.g. RFC 1483
LAN	VoIP	Encapsulation e.g. RFC 1490	Encapsulation e.g. RFC 1483	LAN Switching

FIG. 3 further illustrates a local switch module 206 such as that included in the switch apparatus 102 shown in FIG. 2. As shown in FIG. 3, local switch module 206 includes shared buses 306-A and 306-B that are coupled to NMCs for communicating ATM cells. Local switch module 206 further includes an ingress module 302 and an egress module 304 that are responsible for interfacing ATM cells between the low-speed NMCs and the high-speed ATM cell switching fabric. Preferably, the ingress and egress modules are separately embodied as FPGAs.

Shared buses 306-A and 306-B are preferably each a Cubit-Pro CellBus from TranSwitch Corp. of Shelton, CT. As will be explained in more detail below, the shared buses can be configured for load sharing mode, wherein both buses are active at the same time, or they can be configured for redundancy mode, wherein only one of the buses is active.

FIG. 4 further illustrates an ingress module 302 such as that included in the local switch module 206 shown in FIG. 3.

-9-

As shown, ingress module 302 provides interfaces between the ATM cell switching fabric and three PHY devices. The first interface is with the 16 bit ingress UTOPIA II bus of ALM 402 (such as a LUC4AU01 ATM Layer UNI Manager from Lucent Technologies, for example) of ATM cell switching fabric 208. The ingress module is slave to ALM 402 and responds to 16 PHY addresses coming from the ALM. It runs with a 50 MHz clock 450 generated by the ALM and operates in the cell by cell mode. Cells are 27 words long.

The second interface is with the 8 bit UTOPIA bus in which the ingress module is master to the SAR chip 422 (such as an L64364 ATMizer II+ from LSI, for example) TX Utopia. It runs with the 50 MHz clock generated by the ALM and operates in the cell by cell mode. Cells are 53 bytes long. The SAR (ATMizer) chip 422 is used for segmentation and reassembly of AAL5 frames (used for, e.g., intercard communication, SVC, etc.)

The third and fourth interfaces are with the 16 bit UTOPIA-like inlet buses of Cubit chips 424-A and 424-B (such as TXC-05802 Cubit-Pro Cell Bus switches from TranSwitch Corp. of Shelton, CT) are slaves to the ingress module 302. These buses run with a 40 MHz clock 452 and operate in cell by cell mode. Cells are 28 words long including a 1-word routing tag appended in front of them.

As shown, further included is a CPU interface 426 for allowing configuring and error handling of ingress module 302 by an external CPU.

In this example of the invention, the ingress module 302 emulates 16 PHY devices associated with one port of the ATM cell switching fabric

-10-

208. According to the setup of these 16 UTOPIA addresses, incoming cells are mapped to the three different data paths. Corresponding to each data path, the ingress module 302 maintains three external synchronous FIFOs 404, 406, 408 for cell storage. Each FIFO can store 8 cells. They also provide clock synchronization for the data. The ingress module 302 generates the control signals for these FIFOs.

In general, cells are received from ALM 402 and stored in the proper FIFOs depending on their UTOPIA addresses and the running mode of ingress module 302. The ingress module 302 can be configured for load sharing, in which both Cubit chips 424-A and 424-B are active, or redundancy, where only one of the Cubit chips is set as the active Cubit chip. If one or more of the FIFOs 404, 406, 408 are almost full, the ingress module will not allow the ALM chip to send more cells to those specific FIFOs. This will be achieved by deasserting cell space available signals when responding to ALM's Utopia addresses.

Ingress module 302 further periodically polls SAR (ATMizer) chip 422 and the Cubit chips 424-A and 424-B to see if they can receive a cell. If they can, one cell will be read from the corresponding FIFO and sent to the proper PHY chip. This will be done concurrently for each chip.

For cells going to SAR (ATMizer) chip 422, 27 word cells are converted into 53 byte cells by removing the UDF2 byte of the cells. For cells going to Cubit chips 424-A, 424-B, the first word of the cell coming from ALM 402 is sampled and a Routing Tag is generated out of the first word. The last 4 bits of the Tag are generated from the first 12 bits of the generated Tag with CRC-4 calculation over it.

-11-

As shown in FIG. 4, ingress module 302 includes a data path module 410, an ALM TX Utopia controller module 418, an ATZ FIFO controller module 412, Cubit FIFO controller modules 414 and 416, and a reset and CPU interface module 420.

Data path module 410 receives cells from the FIFOs and performs the word to byte conversion for the SAR (ATMizer) chip 422 and the TAG generation and insertion for the Cubit chips 424-A, 424-B. Control signals for the muxes come from the respective FIFO controller modules 412, 414, 416.

The ALM TX Utopia controller module 418 interfaces with the ALM's ingress UTOPIA bus control signals portion. It also receives cell space available signals from the FIFO control modules. During normal operation, ALM 402 continuously puts the addresses of the 16 PHY devices on the Utopia address bus. ALM TX Utopia controller 418 responds to this polling by putting the cell space available signals coming from the FIFO control modules for each PHY device. These signals are first synchronized to the ALM's 50 MHz clock 450. If ALM 402 starts to send a cell to one of these PHY devices, ALM TX Utopia controller module 418 will first check an SOC alignment. If SOC is not aligned with the cell, it will be ignored. If the cell is normal, it will be written into the FIFO corresponding to the UTOPIA address and the existing setup, as described below.

If load sharing is enabled, cells with a PHY device number 0-7 will be written into Cubit Pro A FIFO 406, and cells with a PHY device number 9-15 will be written into Cubit Pro B FIFO 408. If load sharing is disabled and if Cubit Pro 424-A is active, cells with a PHY device number 0-7 and 9-15 will be written into Cubit Pro A FIFO 406. If load sharing is disabled

-12-

and if Cubit Pro 424-B is active, cells with a PHY device number 0-7 and 9-15 will be written into Cubit Pro B FIFO 408. Cells with PHY device number 8 will always be written into ATZ FIFO 404. After the cell is written into the proper FIFO, one increment cell counter signal will be sent to the proper FIFO controller module.

ATZ FIFO controller module 412 runs at the ALM's 50 MHz clock and keeps the cell counter for the ATZ FIFO 404 and provides interfaces to the ALM Utopia bus. If the SAR (ATMizer) chip 422 has cell space available and the corresponding FIFO cell counter indicates that there is a cell available in FIFO 404, this module starts to read the cell from FIFO 404. During this process it will also generate the control signals for the data path module 410 for muncing and UDF2 byte removing. ATZ FIFO controller module 412 assumes that cells were written into the FIFO with proper SOC alignment (this is the responsibility of ALM TX Utopia controller module 418). But if something goes wrong and a cell comes out of the FIFO with a wrong alignment, this module will generate an SOC Error signal to the Reset Module 420. It will also generate a cell space available signal to the ALM TX Utopia module 418 when appropriate.

Cubit FIFO controller modules 414 and 416 respectively manage Cubit Pro chips 424-A and 424-B running at 40 MHz. They keep the cell counters for the FIFOs and provide interfaces to the Cubit 424-A and Cubit 424-B 16-bit cell inlet buses.

If the Cubit chips 424-A, 424-B have cell space available and the corresponding FIFO cell counter indicates that there is a cell available in one of FIFOs 406, 408, these modules start to read the cell from the corresponding FIFO. During this process they will also generate the control

-13-

signals for the data path module 410 for muxing and tag insertion. The Cubit FIFO controller modules 414, 416 assume that cells were written into the corresponding FIFO 406, 408 with proper SOC alignment (this is the responsibility of ALM TX Utopia controller module 418). But if something goes wrong and a cell comes out of the FIFO with a wrong alignment, these modules will generate SOC Error signals to Reset module 420.

Reset and CPU interface module 420 generates a reset sequence for the following conditions:

(1) Missing SOC signal coming from ATMizer or Cubit FIFOs for cells going from the FIFOs to the PHY devices. This is an error condition on FIFO control, not a missing SOC on the ALM Utopia side to FIFO, in which case the cell is discarded.

(2) ATZ or Cubit chip FIFO overrun condition detected. Normally, the FIFOs can not overflow. When they are almost full, the ALM chip will not be allowed to send more cells using the Utopia control signals.

(3) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

During the reset sequence, the Utopia interfaces are blocked by not allowing any cell operation. Internal logic is reset and the external FIFOs are flushed by generating proper reset signals for them.

Reset and CPU interface module 420 also provides the following pins for access by an external CPU (via CPU interface 426):

(1) Main_reset_ is an asynchronous active low input signal. It has to be asserted for a minimum period of 1 ms. This will initiate a

-14-

reset sequence, which resets the chip, indirectly resets the FIFOs and stops the Utopia interfaces. After main reset is de-asserted, it will take 32 clocks to begin operational mode. At the end of the reset sequence, pin "Err_stat" will be set as an acknowledgment, which signal needs to be cleared using the pin "Clear err cond".

(2) "Err stat" is an asynchronous active high output signal. It will be latched for the following conditions.

(a) Missing SOC signal for cells coming from ATMizer or Cubit chip FIFOs to the PHY devices. This is an error condition on FIFO control. This is not a missing SOC on cells coming from the ALM Utopia side to FIFO, in which case the cells will be discarded and the hardware will realign itself to the in-coming cells by searching for cells with a proper SOC indicator.

(b) FIFO overrun condition detected. Normally the FIFOs can not overflow. When they are almost full, the ALM chip will not be allowed to send more cells to them.

(c) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

When these conditions (a), (b) and (c) are detected, an auto initialization reset sequence is initiated. This will reset the chip and the FIFOs. This will also hold the Utopia interfaces in the blocking state, which will not allow any cell transfer until the reset sequence is completed.

This signal can be cleared with the signal "clear_err_cond".

(3) "Clear_err_cond" is an asynchronous input signal. It will clear the "Err stat" output pin.

-15-

(4) "Share" signal is an asynchronous active high input signal.

0 = Sharing is disabled (default)

1 = Sharing is enabled. In this case, UTOPIA addresses indicating PHY device numbers 0-7 will cause corresponding cells to be directed to Cubit pro 424-A, and UTOPIA addresses indicating PHY device numbers 9-15 will cause cells to be directed to Cubit pro 424-B

(5) "Act_cbt" is an asynchronous input signal. It shows which Cubit pro is the active chip. If sharing is disabled it will have the following functionality:

0 = Cubit pro 424-A is active: UTOPIA addresses indicating PHY device number 0-7 and 9-15 will cause cells to be directed to Cubit pro 424-A

1 = Cubit pro 424-B is active: UTOPIA addresses indicating PHY device number 0-7 and 9-15 will cause cells to be directed to Cubit pro 424-B.

This signal is ignored if sharing is enabled.

After power up, reset and CPU interface module 420 comes up in the reset mode (since Main reset is an active low signal) and stays in the reset mode until Main reset is de-asserted. Before this signal is de-asserted, the ALM, Cubit and SAR (ATMizer) chips need to be initialized to avoid any initial cell loss.

FIG. 5 further illustrates an egress module 304 such as that included in the local switch module 206 shown in FIG. 3.

The egress module 304 provides interfaces between the ATM cell switching fabric 208 and three PHY devices. The first interface is with the

-16-

16 bit egress UTOPIA II bus of the ALM 502 (such as a LUC4AU01 ATM Layer UNI Manager from Lucent Technologies, for example) of the ATM cell switching fabric 208. The egress module 304 is slave to the ALM and responds to single PHY address coming from the ALM 502. It runs with a 50 MHz clock 550 generated by the ALM and operates in the cell by cell mode. Cells are 27 words long.

The second interface with the 8 bit Utopia bus in which the egress module is the master to the SAR chip 526 (such as an L64364 ATMizer II+ from LSI, for example) RX Utopia. It runs with the same 50 MHz clock 550 generated by the ALM and operates in the cell by cell mode. Cells are 53 bytes long. The SAR (ATMizer) chip 526 is used for segmentation and reassembly of AAL5 frames (used for, e.g., intercard communication, SVC, etc.).

The third and fourth interfaces are with the 16 bit Utopia-like Cubit-pro outlet buses of Cubit chips 524-A and 524-B (such as TXC-05802 Cubit-Pro Cell Bus switches from TranSwitch Corp. of Shelton, CT) are slaves to the egress module 304. These buses run with a 40 MHz clock and operate in cell by cell mode. Cells are 27 words long with no Routing tag.

As shown, further provided is a CPU interface 516 for allowing an external CPU access for configuring and error handling of egress module 304.

In this example of the invention, egress module 304 emulates a single PHY device associated with one port of the ATM cell switching fabric 208. Accordingly, the three PHY devices are mapped to a single PHY address.

-17-

Egress module 304 uses three external synchronous FIFOs 518, 520, 522 for cell storage corresponding to each of the three data paths. The ATZ FIFO 518 can store 8 cells. Cubit FIFOs 520, 522 can store 8 or 144 cells depending on the configuration. The FIFOs also provide clock synchronization for the data. The egress module 304 generates the control signals for the FIFOs.

In general, cells are received from ATMizer or Cubit chips and stored in the proper FIFOs. If one or more of the FIFOs 518, 520, 522 are almost full, the egress module 304 will not allow the PHY chips to send more cells to those specific FIFOs. Cells coming from SAR (ATMizer) chip 526 are converted to a 16-bit format before they are written into FIFO 518. During this process, SOC alignment is also checked. If the SOC alignment is not right, the cell will be discarded and a SOC search will start until proper SOC alignment is found from the PHY chips.

Egress module 304 further periodically polls the ATZ and the Cubit chip FIFOs to see if they have a cell available to send to the ALM 502. This is done in a round-robin fashion. If they have a cell to send, the egress module 302 reads the first three words of the cell from the corresponding FIFO. Then it will start to respond to the UTOPIA addresses coming from the ALM as cells are available (only to the address zero). Cell transfer starts as soon as ALM 502 enables the cell read. At the same time, the rest of the cell from the FIFO is clocked in.

As shown in FIG. 5, egress module 304 includes a data path module 504, an ALM RX Utopia controller module 514, an ATZ FIFO controller

-18-

module 506, Cubit FIFO controller modules 508, 510, and a reset and CPU interface module 512.

Data path module 504 receives cells from the FIFOs 518, 520, 522 and performs cell multiplexing for the ALM 502. Selection signals come from the ALM RX Utopia controller module 514.

The ALM RX Utopia controller module 514 interfaces with the ALM's egress UTOPIA bus control signals portion. It also receives cell space available signals from the FIFO control modules 506, 508, 510. During normal operation, ALM 502 continuously polls the addresses of the single PHY device on the Utopia address bus. ALM RX Utopia controller 514 also continuously polls the FIFO cell availability signals coming from FIFO controller modules 506, 508, 510. If they indicate that there is a cell available, the corresponding PHY chip will be scheduled for the next transfer. As soon as the current cell transfer operation is completed, the first three words of the next cell will be read from the FIFO. After this, ALM RX Utopia controller 514 starts to respond to the ALM's polling by issuing cell space available signals only for the device which was previously scheduled.

If ALM 502 asserts the read enable signal, ALM RX Utopia controller module 514 will start to send the cell to the Utopia bus. At the same time it will start to read the rest of the cell which was waiting in the FIFO. During this process, the ALM RX Utopia controller module 514 assumes that cells were written into the FIFO with proper SOC alignment (this is the responsibility of FIFO controller module 506, 508, 510). But if something goes wrong and a cell comes out of the FIFO with a wrong alignment, this module will

-19-

generate an SOC Error signal to the Reset module. The ALM RX Utopia controller module 514 responds to single UTOPIA addresses of zero.

ATZ FIFO controller module 506 runs at the ALM's 50 MHz clock 550 and keeps the cell counter for the ATZ FIFO and provides interfaces to the ALM Utopia bus. If SAR (ATMizer) chip 526 has a cell to send and the corresponding FIFO cell counter indicates that there is room in FIFO 518, this module starts to read the cell from SAR (ATMizer) chip 526. During this process, it checks the SOC alignment. If SOC is not aligned with the cell, it will discard the cell and start searching for properly SOC aligned cells. If SOC is aligned properly, this module will do the byte to word conversion by using external staging registers and will write the cells in to the FIFO 518. It will also generate a cell available signal to the ALM RX Utopia controller module 514.

Cubit A and B FIFO controller modules 508, 510 respectively manage Cubit pro 524-A and 524-B chips running at 40 MHz clock 552. They keep the cell counters for the FIFOs and provide interfaces to the Cubit 524-A and Cubit 524-B 16-bit cell outlet buses.

If one of the Cubit chips 524-A and 524-B has a cell available and the FIFO cell counter indicates that there is room available in the respective FIFO 520, 522, these modules will start to read the cell from the corresponding Cubit chip. During this process they will check SOC alignment. If SOC is not aligned with the cell, they will discard the cell and start searching for properly SOC aligned cells. If SOC is aligned properly, they will write the cells into the corresponding FIFO. They will also

-20-

generate a cell available signal to the ALM RX Utopia controller module 514.

Reset and CPU interface module 512 generates a reset sequence for the following conditions:

(1) Missing SOC signal coming from ATMizer or Cubit FIFOs for cells going to data path module 504. This is an error condition on FIFO control, in which case the unaligned cell will be discarded.

(2) ATZ or Cubit chip FIFO overrun condition detected.

Normally, the FIFOs can not overflow. When they are almost full, PHY chips will not be allowed to send more cells to them using the Utopia control signals.

(3) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

During the reset sequence the Utopia interfaces are blocked by not allowing any cell operation. Internal logic is reset and the external FIFOs are flushed by generating proper reset signals for them.

Reset and CPU interface module 512 also provides the following pins for access by an external CPU (via CPU interface 516):

(1) "Main_reset_" is an asynchronous active low input signal. It has to be asserted for a minimum period of 1 ms. This will initiate a reset sequence, which resets the chip, indirectly resets the FIFOs and stops the Utopia interfaces. After main reset is de-asserted, it will take 32 clocks to go to operational mode. At the end of the sequence, pin "Err_stat" will be set as an acknowledgment, which signal needs to be cleared using the pin "Clear_err_cond".

-21-

(2) "Err_stat" is an active high output signal. It will be latched for the following conditions.

(a) Missing SOC signal for cells coming from the ATMizer or Cubit chip FIFOs to the egress module. This is an error condition on FIFO control, not a missing SOC on the ALM Utopia side to FIFO, in which case the unaligned cell will be discarded and hardware will realign itself to the in-coming cells by searching for cells with a proper SOC indicator.

(b) FIFO overrun condition detected. Normally the FIFOs should not overflow, and so this can only be a controller error. When the FIFOs are almost full, PHY chips will not be allowed to send more cells to them.

(c) After each Main reset is given to the chip, the same reset sequence will be started and this condition will be latched on the Err_stat pin as well.

When these conditions (a), (b) and (c) are detected, an auto initialization reset sequence is initiated. This will reset the chip and the FIFOs. This will also hold the Utopia interfaces in the blocking state, in which they will not be allowed to perform any cell transfers until the reset sequence is completed.

This signal can be cleared with the signal "clear_err cond".

(3) "Clear_err_cond" is an asynchronous input signal. It will clear the "Err_stat" output pin.

(4) "FIFO_Threshold" is an asynchronous input signal. It will determine the depth of the FIFOs for the Cubit chips. This value has to be set up initially and should not be changed during normal operation.

0 = Threshold is 8 cells

-22-

1 = Threshold is 144 cells

After power up, this module comes up in the reset mode (since Main reset is an active low signal) and stays in the reset mode until Main reset is deasserted. Before this signal is de-asserted, the ALM, Cubit and ATMizer chips need to be initialized to avoid any initial cell loss.

FIG. 6 further illustrates a narrowband card 202, 204 that can be included in the switch apparatus 102 illustrated in FIG. 2. It includes a cell bus controller 602 that communicates with two Cubit chips 628-A and 628-B, and a virtual circuit (VC) controller 608. The VC controller 608 further communicates with a packet controller 618, an ATM cell controller 620, and a voice/fax controller 626 via a shared bus 616. The cell bus controller 602 further communicates with first AAL1 SAR chip 604 and second AAL1 SAR chip 606.

The cell bus controller 602 is preferably implemented as a FPGA and provides five Utopia interfaces - - between the two Cubit chips 628-A and 628-B, the first AAL1 SAR chip 604, the second AAL1 SAR chip 606 and the VC controller 608. The cell bus controller 602 plays the role of an ATM layer multiplexer device providing interfaces to the shared buses 306-A and 306-B of the local switch module 206 from multiple PHY devices with different priorities. That is, the cell bus controller 602 multiplexes ATM cells from the low-speed interfaces connected to the NMC with the high-speed port of the ATM cell switching fabric, which the NMC shares with other NMCs via the shared bus of the local switch module 206.

The first and second AAL1 SAR chips 604, 606 are, for example, PMC73121 AALigator II chips from PMC-Sierra and are programmed to be in the single PHY mode. Between the SAR chips and the cell bus controller

-23-

there are FIFOs (not shown) which isolate the clocking domains. The SAR's Utopia runs at 33MHz. The FIFOs are, for example, SuperSync device IDT72261 from IDT of Santa Clara, CA. The first and second AAL1 SAR chips 604, 606 provide interfaces to ATM networks using T1/E1, T3/E3, and/or supported services.

VC controller 608 is preferably implemented by, for example, an L64364 ATMizer II+ from LSI. The VC controller's Utopia port is also configured to be in the single PHY mode. The Utopia clock runs at 40 MHz which is synchronous to the Cubit's Utopia clock. Also, the ATMizer is preferably configured to ignore parity on the Utopia bus.

As shown, VC controller 608 implemented by, for example an ATMizer chip, includes a SAR AAL0/AAL5 engine 610, a plurality of VC queues 612, and a multi-service engine 614.

Generally, ATM cells received from, or to be sent to, the narrowband interfaces via shared bus 616 are stored in the VC queues 612. The rates at which the VC queues 612 are respectively serviced are determined by the quality of service for the particular flows to which the ATM cells stored in the VC queues belong. SAR engine 610 performs segmentation and reassembly functions for AAL5 frames and cell forwarding for AAL0 frames. It also runs a schedule-based service algorithm to determine which VC queue should be serviced (i.e., for transmitting a cell) at each cell time. VC queues 612 are implemented by a linked list of buffers. Multi-service engine 614 is a software module that performs packet translation if necessary for adapting flows for ATM cell switching in the switching fabric (e.g. FR to ATM network or service interworking), service functions based on header lookup, flow to VC mapping, and queuing of cells to the

-24-

appropriate per-VC queue 612. An example of a method and apparatus that can be used to implement the management of quality of service requirements for various flows according to the invention is described in co-pending U.S. Provisional Appln. No. 60/_____ Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference.

Shared bus 616 is preferably a PCI bus adapted for transfers of 1 Gbps.

Packet controller 618 interfaces between packets and frames that are input/output via attached IP/Frame Relay networks and ATM cells that are input/output via cell bus controller 602 and converted to packets and frames by VC controller 608. Preferably, it includes a HDLC controller (such as a PEB20324 from Siemens, for example) which performs HDLC functions such as bit stuffing/ unstuffing, CRC checking, etc. Packets and frames received by packet controller 618 from attached IP / Frame Relay networks are processed by the HDLC controller and sent to multi-service engine 614 via shared bus 616 for conversion to ATM cells and queuing in per-VC queues 612. Conversely, ATM cells destined for attached IP / Frame Relay networks are converted to packets by multi-service engine 614 and forwarded to packet controller 618 via shared bus 616, which then immediately forwards them to the appropriate IP / Frame Relay network port.

ATM cell controller 620 forwards ATM cells that are input from attached ATM networks to multi-service engine 614 via shared bus 616 and forwards ATM cells destined to attached ATM networks that are received

-25-

from the switch fabric via cell bus controller 602 to the attached ATM networks.

As shown, ATM cell controller 620 preferably includes an ATM service engine 622 and an ATM PLCP 624. ATM service engine 622 performs dual leaky bucket UPC based on ATM Forum 4.0 and uses a VPI/VCI index into a table to find the corresponding VC queue when forwarding cells between multiservice engine 614 and attached ATM networks. ATM PLCP 624 performs ATM physical layer functions according to ITU-.432 and direct cell mapping to DS1 or E1 transmission systems according to ITU-T G.804.

Voice/fax controller 626 converts voice/fax data that is received via attached PSTN networks into packets that are then forwarded to multi-service engine 614 via shared bus 616, and likewise converts packetized voice/data that is destined for attached PSTN networks and is received from multi-service engine 614 via shared bus 616 into voice/fax data for forwarding over the attached PSTN networks. An example of a method and apparatus that can be used to implement the management of voice/fax flows according to the invention is described in co-pending U.S. Provisional Appln. No. 60/_____, Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference.

FIG. 7 further illustrates a voice/fax controller 626 that can be included in the narrowband line card 202, 204 illustrated in FIG. 6. As shown, it preferably includes a digital signal processor (DSP) service engine 702, a DSP controller 706, and a set of DSPs 708-1...708-N. DSP

-26-

service engine 702 is preferably implemented by a LSI L64364 ATMizer II+. DSPs 708 are preferably implemented by TMS320C6201 chips.

As shown, DSP controller 706 provides the interface between the DSP service engine and the DSPs. DSP service engine 702 communicates with DSP controller 706 via a UTOPIA data path interface from which it receives, decodes and executes messages. DSP service engine 702 also controls the DSP controller 706 via its secondary port interface (described below).

DSP service engine 702 operates in slave mode, with cell data outputs from its transmit FIFO (not shown) going to the UTOPIA master, and cell data inputs from DSP controller 706 master being input to its receive FIFO 704. DSP controller 706 accesses the DSPs via its host port interface (described below). DSP service engine is mainly responsible for performing media type adaptation of voice/fax flows received or sent over PSTN networks connected to switch apparatus 102.

A DSP service engine 702 that can be adapted for use in the present invention is described in co-pending U.S. Provisional Application No. _____, Atty. Dkt. #85323/239638, commonly owned by the assignee of the present invention, the contents of which are incorporated herein by reference. Briefly, however, TDM-based voice/fax data streams received by voice/fax controller 626 from attached PSTN networks (via a time slot interchanger known in the art, for example) are packetized by the DSPs and stored in DSP output queues (not shown). DSP controller 706 forwards the voice/fax packets from the DSP output queues to DSP service engine 702. DSP service engine 702 then assembles the packets in accordance with the egress media type for the voice/fax connection. If the

-27-

egress media type is IP, for example, the DSP service engine adds a RTP, UDP, and IP header to the packets and forwards them to multi-service engine 614 via shared bus 616. Conversely, packetized voice/fax data streams are received from multi-service engine 614 via shared bus 616. DSP service engine 702 then converts the packets of the ingress media type into voice/ fax data. For example, if the ingress media type was IP, DSP service engine 702 reads the RTP, UDP and IP headers of the packets and writes the packets to the appropriate DSP input queues (not shown). The DSPs convert the packetized data into voice/fax streams for output via attached PSTN networks.

FIG. 8 further illustrates a DSP controller 706 such as that illustrated in FIG. 7. As shown, it includes a cell receive block 802, a cell transmit block 804, a command processor 808, a secondary port interface (SPI) block 806, a host port interface (HPI) block 814, and block read unit 812, and a block write unit 810.

The cell receive block 802 interfaces with the DSP service engine 702 receive port. DSP controller 706 acts as Utopia master and DSP service engine 702 as slave. The cell receive block 802 is responsible for loading the messages into internal RAM buffers (not shown), which provide storage for two cells. Command processor 808 issues a load command to cell receive block 802, and cell receive block 802 generates a ready signal when the corresponding buffer contains valid data. The command processor 808 then decodes and executes the contents of the current buffer while cell receive block 802 is loading the next buffer.

Cell receive block 802 checks the parity of the receive data bus and compares it to the incoming parity bit. The parity error bit in Interrupt 0

-28-

vector register is set and the DSP controller 706 asserts an interrupt signal when a parity error is detected.

Cell receive block 802 saves the parity bits for the header word, which bits are looped back for a Block Read operation when the header word is written to the external FIFO.

The cell buffers (not shown) are implemented using embedded 256 by 8 bit wide RAMS. The command processor 808 is able to overlap the reading of the buffers with transferring of data to the host port interface 814 by using a "prefetch" mechanism. The command processor 808 asserts a next word load signal and cell receive block 802 automatically reads the next 4 bytes and form a 32 bit word at the buffer data outputs.

Cell transmit block 804 interfaces to DSP service engine 702's transmit port. The block read unit 812 directly writes cell data to external FIFO 704 while the almost full flag remains de-asserted.

Immediately after reset, cell transmit block 804 waits until the block read unit 812 has programmed the external FIFO almost full and almost empty flags. When the external FIFO 704 contains at least one full cell, the almost empty flag will be de-asserted. When the DSP service engine 702 asserts a signal indicating that a cell is available the cell transmit block 804 will proceed to generate the control signals to read the data out of the external FIFO 704 and into the DSP service engine 702. The cell transmit block 804 maintains an octet counter to keep track of the words read out of the external FIFO and to generate a start of cell signal.

A read enable signal is input to the external FIFO 704, which signal is delayed one clock, thus indicating valid data on the input pins. The external FIFO almost empty flag is valid on the second clock

-29-

following a read cycle and it is registered at the input of the DSP controller 706. The cell transmit state machine 804 samples the internal almost empty flag three clock cycles after the last byte of a cell have been read out of the external FIFO to accommodate the latencies.

SPI block 806 communicates with the secondary port interface of the DSP service engine 702 so as to provide direct access to DSP memory locations, HPI block registers, DSP reset control, reading and clearing of the Interrupt 0 vector register, and reading of the Interrupt 1 vector register.

The SPI runs asynchronously to the DSP controller 706. SPI block 806 detects the falling edge of an access signal from the SPI to initiate a secondary port access to the DSP controller 702, and asserts a ready signal when the cycle is completed. Two address lines and a write enable signal are used to provide access to four addressable locations within the DSP controller 702: a read/write control register, a read/write data register, the Interrupt 0 vector register, and the Interrupt 1 vector register.

The read/write control register is loaded first and is used to set up the starting addresses for reads and writes, as well as selecting one of the DSPs. The OpCode field of the read/write control register and the write enable signal determine the behavior of the read/write operation. An OpCode value of "100" causes a read or write operation to the selected address of the selected DSP to be performed. An OpCode value of "101" causes a write operation into the DSP Reset register. OpCode values of "000" through "011" cause the indicated HPI register of the selected DSP to be accessed for diagnostic purposes.

-30-

Selection of the Interrupt 0 and 1 vector registers allows the contents of these registers to be viewed.

Command processor 808 issues buffer load requests to cell receive block 802. When the cell receive block 802 has loaded 64 bytes into a cell buffer it will assert the corresponding buffer ready flag. The command processor 808 decodes the message type field of the header word and asserts a start signal to the block write unit 810 or the block read unit 812. The command processor 808 issues a buffer load request to the alternate buffer while the block write or block read units 810, 812 are operating on the data from the first cell buffer. The command processor 808 also muxes the data and control signals between the cell

When block write unit 810 receives a start signal from the command processor 808, it loads the cell data message length field into a message counter register, loads the DSP ID field into a register, and asserts a signal to "prefetch" the read/write DSP address word from the buffer. It asserts an initial request signal to the HPI to let it know that this transfer will load the address register as well as perform one 32 bit data write cycle to the corresponding DSP.

Block write unit 810 waits for an acknowledge signal to be asserted by the HPI. When the HPI acknowledges that the DSP's HPIA register has been loaded, the block write unit 810 will "prefetch" the first data word to be written and again wait for the acknowledge signal. For subsequent write cycles the block write unit 810 asserts subsequent request signals.

Block write unit 810 uses the message counter to keep track of how many 4 byte words are to be written, and it also uses a word counter

-31-

to keep track of the number of 4 byte words remaining in this cell. It then determines if another cell is needed to complete the block write operation. When all the words in the current cell have been written, a signal is asserted indicating this. It should be noted that the first cell has a payload of 14 words (56 bytes) and subsequent cells have a payload of 15 words (60 bytes).

When block write unit 810 is finished writing all the words indicated by the message counter, it generates a signal to the SPI to assert the End of Block interrupt, which signal further identifies which bit to assert in the Interrupt 0 vector register.

Block read unit 812 drives the data and control signals to the TX FIFO 704. After reset it generates the necessary control signals to program the almost empty flag to be asserted when less than one cell is in the TX FIFO (64 bytes) and the almost full flag is asserted when less than one cell (64 bytes) can be written. When block read unit 812 receives a start signal from the command processor 808, it loads the cell data message length field into a message counter register, loads the DSP ID field into a register, and stores the entire cell data into the header register.

The block read unit then checks the TX FIFO almost full flag. If there is room for an entire cell it writes the header register and asserts a signal to "prefetch" the read/write DSP address word from the buffer. It asserts the initial request signal to the HPI to let it know that this transfer will load the address register as well as perform one 32 bit data read cycle to the corresponding DSP. If there is no room in TX FIFO 704 for an entire cell, the block read unit 812 "aborts" and the cell is not read.

-32-

Block write unit 810 waits for an acknowledge signal to be asserted by HPI 814. When the HPI acknowledges that the DSP HPIA register has been loaded, and the first 4 byte word has been read, block read unit 812 writes the data one byte at a time to the TX FIFO 704. Block read unit 812 uses the word counter to determine when a new cell is started and when to write the header register into the TX FIFO.

When the number of words to read does not fill an integral number of cells, block read unit 812 starts the FILL state machine. Its function is to take over and "fill" up the cell with data while the block read unit 812 indicates to the command processor 808 that it is done.

Host port interface 814 is responsible for arbitrating between the secondary port requests and the block read or block write unit requests. Priority is given to SP requests; however, when the block read or block write units 810, 812 assert their initial request signals, the address and data portions of the transfers are not interruptible by the SPI 806.

The arbiter will first check for SPI requests. If an SP initial request signal is asserted, the opcode bits are checked to determine if the secondary port will perform a memory read/ write ("100") or if direct access to the DSP HPI registers is required. For direct access the opcode bits will drive the control bits. For memory transfers, the state machine will indicate HPIA during the address cycle and HPID during data cycle. If the secondary port performs a memory access, the DSP ID is checked against the DSP ID of the last block read or block write transfer and the HIT flag is set if the secondary port access is to the same DSP.

-33-

If no SP request is asserted, then the block write and block read request signals are checked. If an initial request is made or the HIT flag is true, then the HPIA address register is loaded and a read or write cycle is performed and the HIT flag is cleared. If a data request transfer is made and the HIT flag is false then there is no need to reload the HPIA register.

In the above description of an example of a narrowband card 202, 204 according to the present invention, egress channel adaptation is simply the reverse process of ingress channel adaptation. That is, if necessary, ATM cells forwarded to the narrowband card via another narrowband card (via the shared bus of local switch module 206) or from a broadband card (via switch fabric 208 and local switch module 206), are converted into packets or frames as appropriate for forwarding on an attached non-ATM network.

It should be noted that broadband card BSC1 210 also preferably includes similar functionality for adapting ingress and egress broadband flows for switching, for example by enqueueing ATM cells communicated or to be communicated over cell switching fabric 208 in virtual circuit queues, and transforming, if necessary, between such ATM cells and any other media type in accordance with the network on the broadband network connections of BSC1. For example, if an egress media type is Ethernet, the BSC1 reassembles the ATM cells received from cell switching fabric 208 into an AAL5 frame, performs service functions, and enqueues the frame to the corresponding Ethernet MAC Tx queue. Conversely, if an ingress media type is Ethernet, the BSC 1 segments

-34-

AAL5 frames from the corresponding Ethernet MAC Rx queue into ATM cells for switching via cell switching fabric 208.

It should be further apparent that the present invention allows for switching flows of data between two broadband networks or two narrowband networks, as well as allowing for switching flows between a broadband network and a narrowband network as is more explicitly described above.

Although the present invention has been described in detail with reference to the preferred embodiments thereof, those skilled in the art will appreciate that various substitutions and modifications can be made to the examples described herein while remaining within the spirit and scope of the invention as defined in the appended claims.

-35-

What is claimed is:

1. A carrier class switch apparatus integrated in a single switching platform comprising:
 - a switching fabric adapted to switch traffic between a plurality of broadband switching ports;
 - a broadband interface coupled to one of the plurality of broadband switching ports; and
 - a local switch module coupled to another one of the plurality of broadband switching ports and to one or more narrowband interfaces.
2. An apparatus according to claim 1, wherein the broadband interface is adapted to be further coupled to one or more broadband networks.
3. An apparatus according to claim 1, wherein the one or more narrowband interfaces are each adapted to be further coupled to one or more narrowband networks.
4. An apparatus according to claim 1, further comprising a switch control card coupled to the switching fabric, the broadband interface and the local switch module, wherein the switch control card is adapted to route and manage virtual circuit connections between the plurality of broadband switching ports in accordance with a quality of service requirement.

-36-

5. An apparatus according to claim 1, wherein the local switch module is adapted to cause the one or more narrowband interfaces to emulate the 5 broadband interface.

6. An apparatus according to claim 1, wherein the local switch module comprises:

a shared bus coupled to the one or more narrowband interfaces;

an ingress module coupled to the shared bus and the switching fabric that is adapted to forward traffic from the switching fabric to appropriate narrowband interface; and

an egress module coupled to the shared bus and the switching fabric that is adapted to forward traffic from the narrowband interfaces to the shared bus.

7. An apparatus according to claim 6, wherein the shared bus is comprised of a plurality of buses configured for load sharing mode.

8. An apparatus according to claim 6, wherein the shared bus is comprised of a plurality of buses configured for redundancy mode.

9. An apparatus according to claim 6, wherein the ingress module includes:

-37-

one or more data paths between the switching fabric and the shared bus, the data paths adapted to buffer traffic between the switching fabric and corresponding narrowband interfaces; and

a controller module adapted to communicate with the switching fabric and to map traffic received from the switching fabric to the one or more data paths.

10. An apparatus according to claim 6, wherein the egress module includes:

one or more data paths between the switching fabric and the shared bus, the data paths adapted to buffer traffic between the switching fabric and corresponding narrowband interfaces; and

a controller module adapted to communicate with the switching fabric and to map traffic received from the narrowband interfaces to the switching fabric.

11. An apparatus according to claim 9, wherein the egress module includes:

one or more data paths between the switching fabric and the shared bus, the data paths adapted to buffer traffic between the switching fabric and corresponding narrowband interfaces; and a controller module adapted to communicate with the switching fabric and to map traffic received from the narrowband interfaces to the switching fabric.

12. An apparatus according to claim 6, wherein at least one of the narrowband interfaces comprises:

-38-

at least one network interface adapted to communicate traffic with a narrowband network;

a virtual circuit queue for buffering traffic associated with the narrowband network; and

a cell bus controller coupled to the virtual circuit queue and the shared bus, the cell bus controller adapted to multiplex the narrowband network traffic buffered in the virtual circuit queue with the broadband switching port of the switching fabric to which the shared bus is coupled.

13. An apparatus according to claim 11, wherein at least one of the narrowband interfaces comprises:

at least one network interface adapted to communicate traffic with a narrowband network;

a virtual circuit queue for buffering traffic associated with the narrowband network; and

a cell bus controller coupled to the virtual circuit queue and the shared bus, the cell bus controller adapted to multiplex the narrowband network traffic buffered in the virtual circuit queue with the broadband switching port of the switching fabric to which the shared bus is coupled.

14. An apparatus according to claim 12, wherein the at least one narrowband interface further includes a multi-service engine that translates between a first media type associated with the narrowband network traffic and a second media type of the switching fabric.

-39-

15. An apparatus according to claim 14, wherein the first media type is IP and the second media type is ATM.
16. An apparatus according to claim 14, wherein the first media type is ATM and the second media type is ATM.
17. An apparatus according to claim 14, wherein the first media type is frame relay and the second media type is ATM.
18. An apparatus according to claim 14, wherein the first media type is voice and the second media type is ATM.
19. An apparatus according to claim 14, wherein the first media type is fax and the second media type is ATM.
20. An apparatus according to claim 2, wherein the broadband interface includes a multi-service engine that translates between a first media type associated with broadband network traffic and a second media type of the switching fabric.
21. An apparatus according to claim 14, wherein the broadband interface is adapted to be further coupled to one or more broadband networks, and the broadband interface includes a second multi-service engine that translates between a third media type associated with broadband network traffic and the second media type of the switching fabric.

-40-

22. A carrier-class switch apparatus integrated in a single switch platform comprising:

a switching fabric adapted to switch ATM cells between a plurality of broadband ports;

a broadband interface adapted to be coupled between one of the broadband ports and a broadband network, the broadband interface including a broadband multi-service engine that is adapted to translate between ATM cells switched via the switching fabric and an egress media type associated with, broadband network traffic;

a local switch module coupled to another of the broadband ports, the local switch module including a shared bus; and

a narrowband interface adapted to be coupled between the shared bus of the local switch module and a narrowband network, the narrowband interface including a narrowband multi-service engine that is adapted to translate between ATM cells switched via the switching fabric and an ingress media type associated with narrowband network traffic.

23. An apparatus according to claim 22, wherein the ingress media type is IP.

24. An apparatus according to claim 22, wherein the ingress media type is ATM.

25. An apparatus according to claim 22, wherein the ingress media type is frame relay.

-41-

26. An apparatus according to claim 22, wherein the ingress media type is voice.

27. An apparatus according to claim 22, wherein the egress media type is IP.

28. An apparatus according to claim 22, wherein the egress media type is ATM.

29. An apparatus according to claim 22, wherein the egress media type is frame relay.

30. An apparatus according to claim 22, wherein the egress media type is voice.

31. An apparatus according to claim 26, wherein the egress media type is IP.

32. An apparatus according to claim 26, wherein the egress media type is ATM.

33. An apparatus according to claim 26, wherein the egress media type is frame relay.

-42-

34. An apparatus according to claim 26, wherein the egress media type is voice.

35. A method for switching traffic between networks in an integrated carrier-class switching platform, the method comprising:

- (a) communicating broadband network traffic via a broadband interface;
- (b) adapting the broadband network traffic for switching via a broadband switching port of a switching fabric;
- (c) communicating narrowband network traffic via at least one narrowband interface;
- (d) multiplexing the narrowband network traffic on a shared bus; and
- (e) adapting the narrowband network traffic for switching via another broadband switching port of the switching fabric so that the at least one narrowband interface emulates the broadband interface.

36. A method according to claim 35, wherein the steps of adapting the broadband and narrowband network traffic for switching each include converting the network traffic from a first media type to a second media type:

37. A method according to claim 36, wherein the second media type is ATM.

-43-

38. A method according to claim 37, wherein the first media type is IP.

39. A method according to claim 37, wherein the first media type is ATM.

40. A method according to claim 37, wherein the first media type is voice.

41. A method according to claim 37, wherein the first media type is frame relay.

42. An apparatus for switching traffic between networks in an integrated carrier-class switching platform, the apparatus comprising:
means for communicating broadband network traffic via a broadband interface;
means for adapting the broadband network traffic for switching via a broadband switching port of a switching fabric;
means for communicating narrowband network traffic via at least one narrowband interface;
means for multiplexing the narrowband network traffic on a shared bus; and
means for adapting the narrowband network traffic for switching via another broadband switching port of the switching fabric so that the at least one narrowband interface emulates the broadband interface.

-44-

43. An apparatus according to claim 42, wherein the means for adapting the broadband and narrowband network traffic for switching each include means for converting the network traffic from a first media type to a second media type.

1/8




FIG. 1





FIG. 2

3/8

**FIG. 3**

4/8

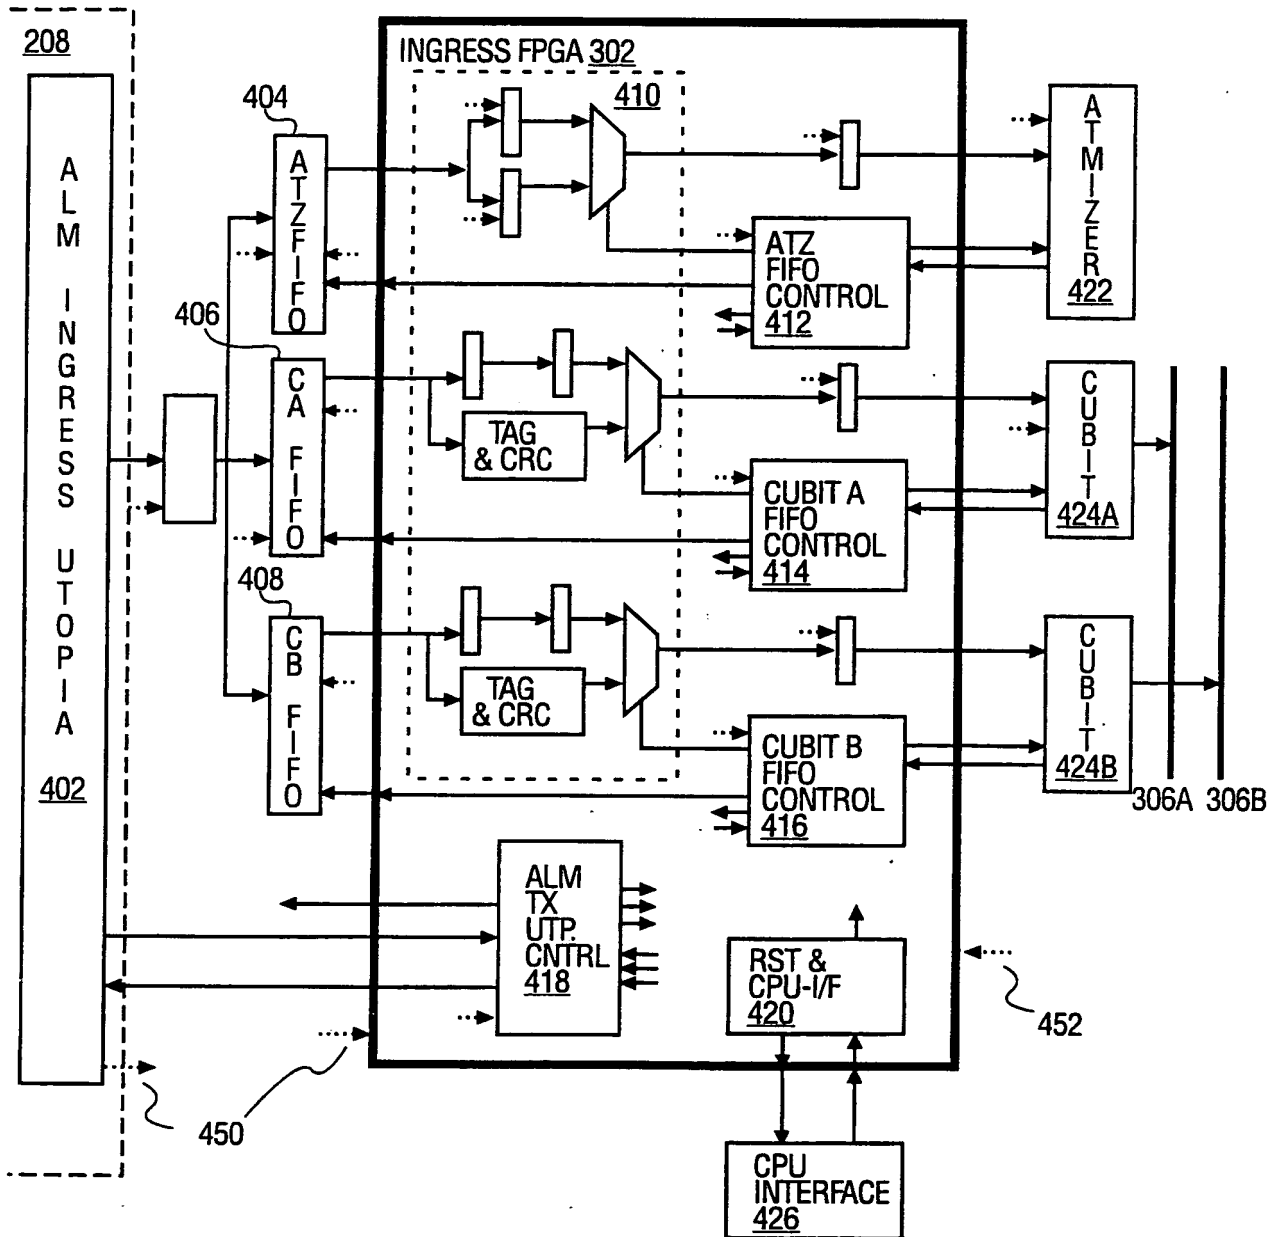


FIG. 4

5/8

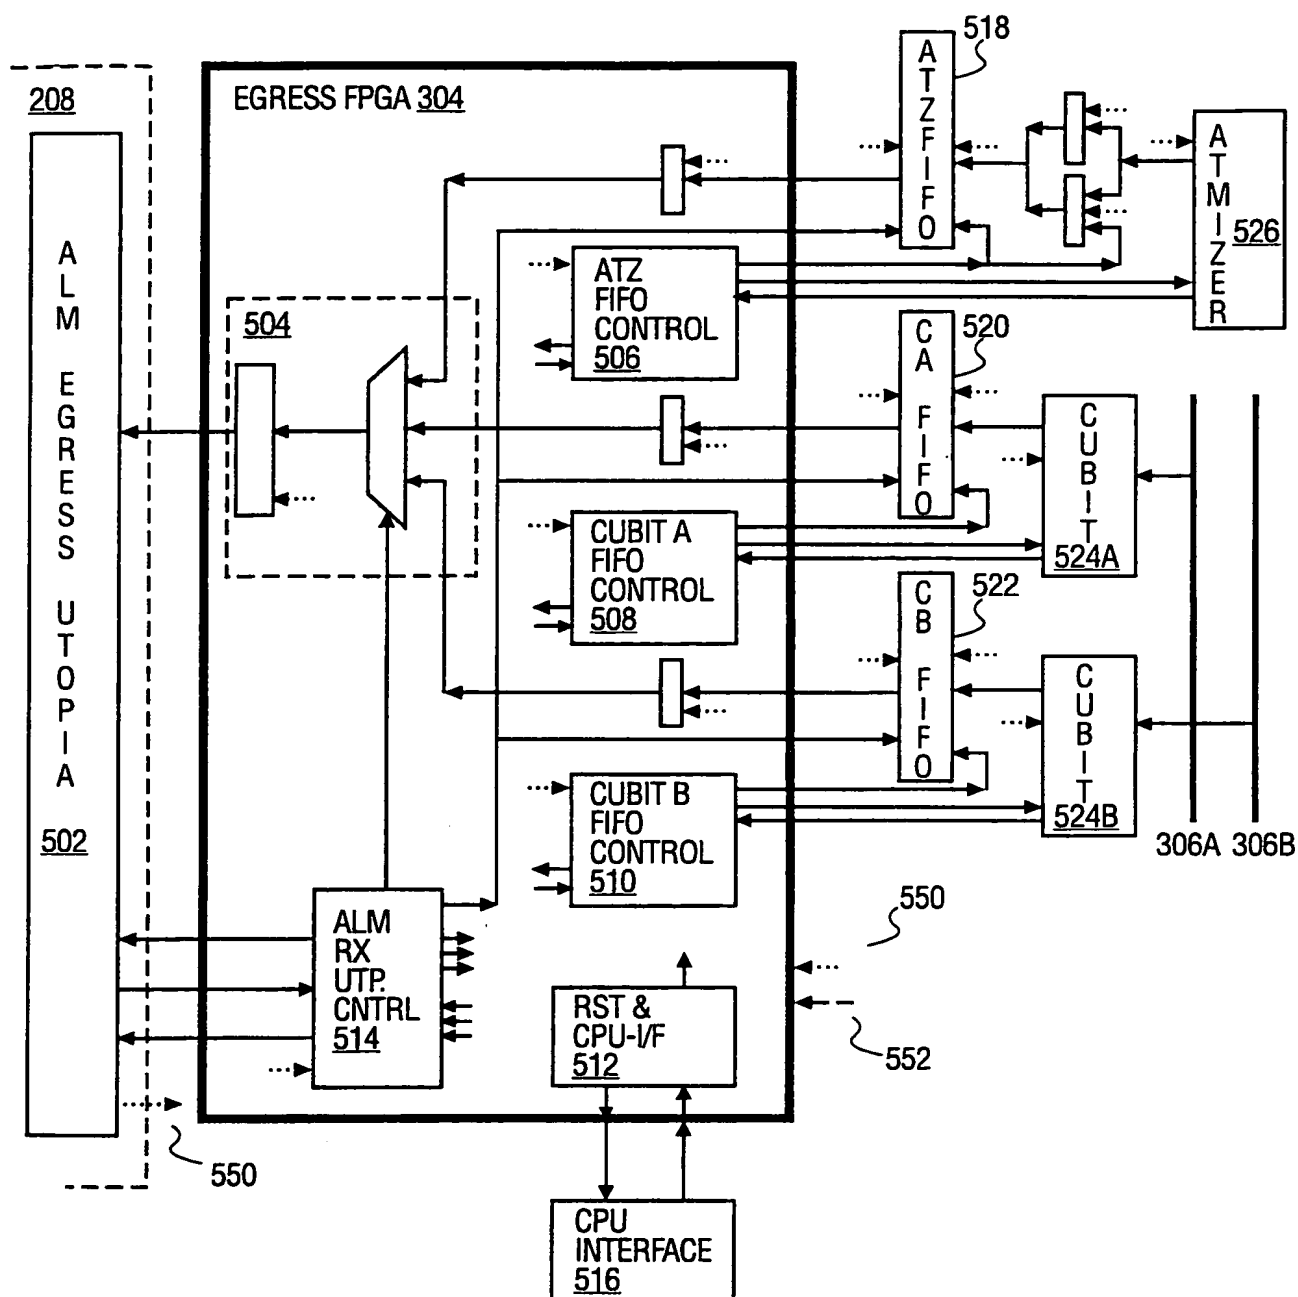


FIG. 5

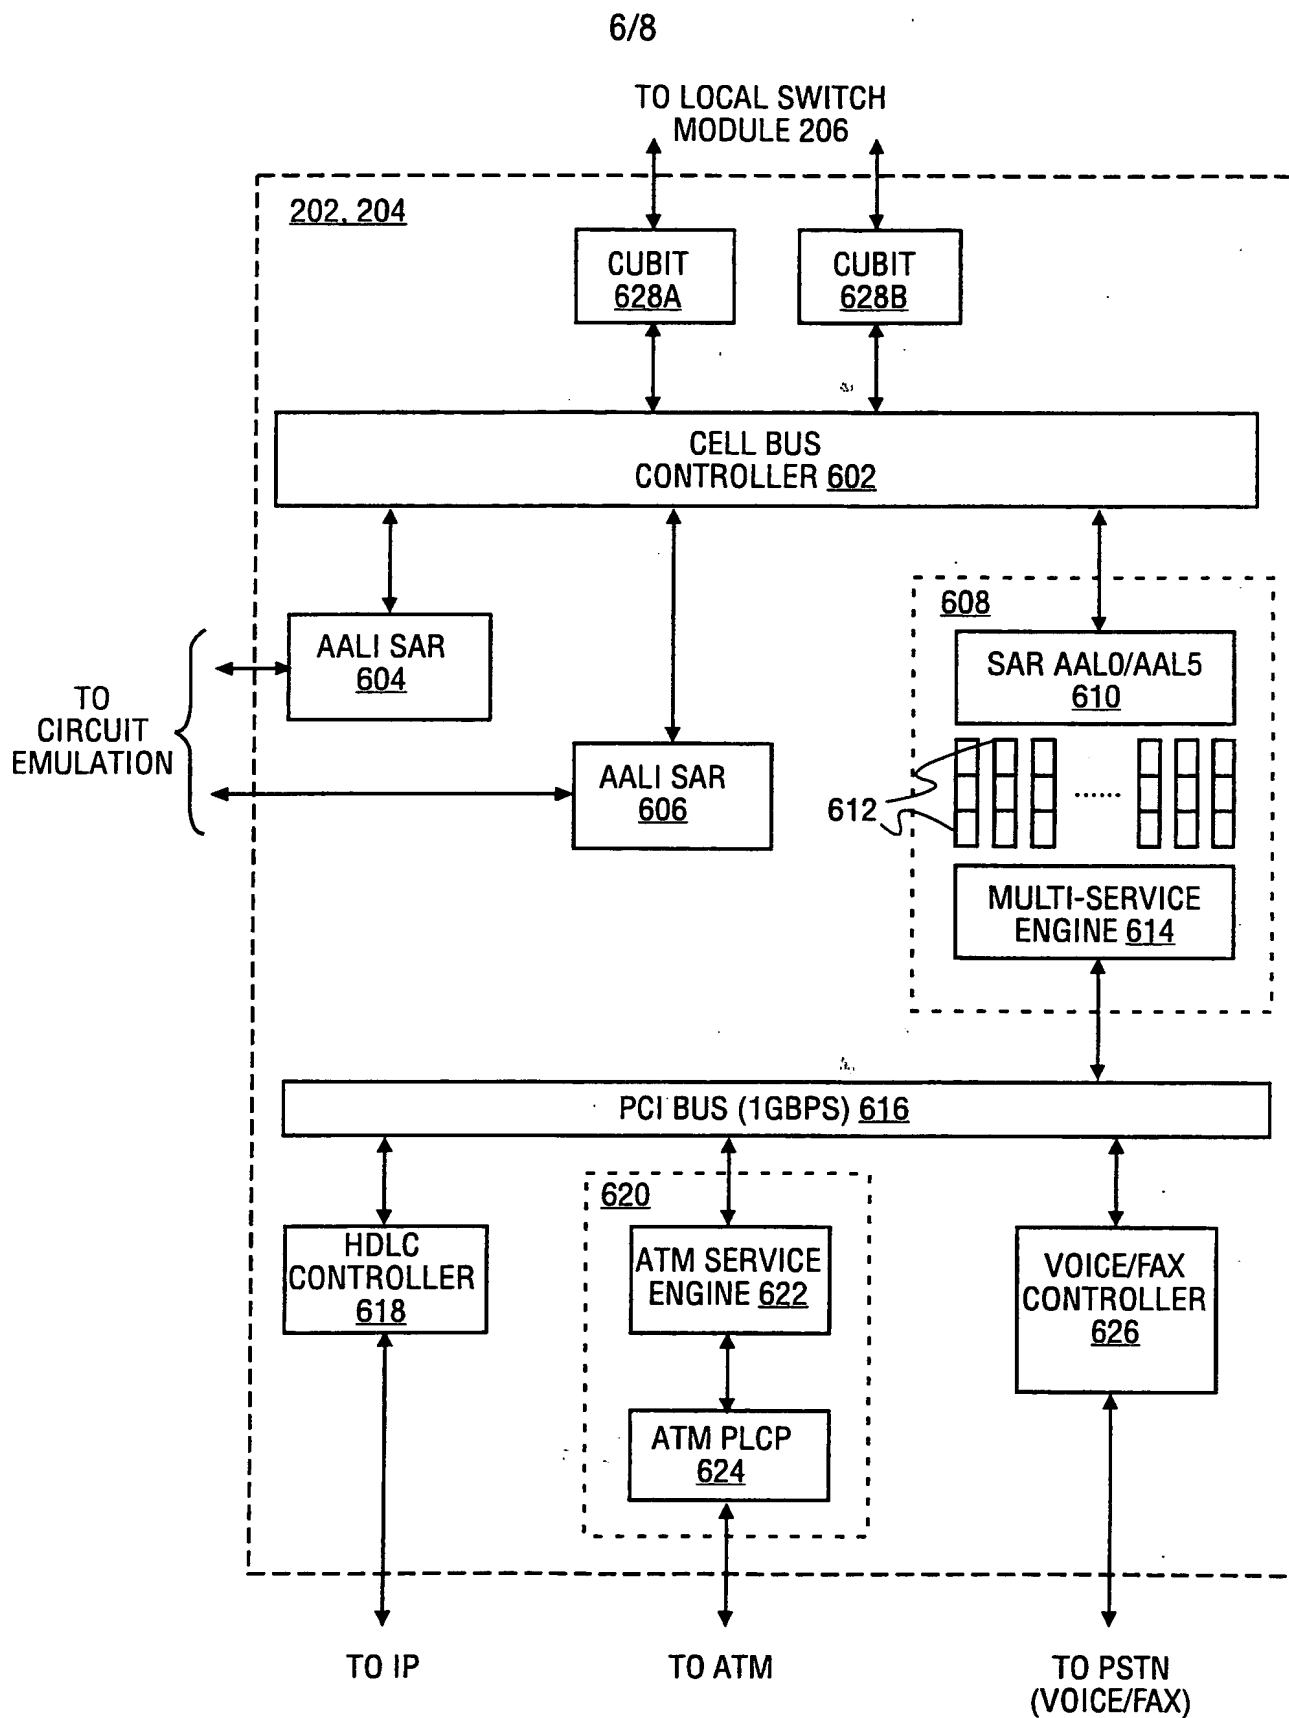


FIG. 6




FIG. 7

8/8




FIG. 8

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 00/06481

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/64 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US 5 809 022 A (RUNYON JAMES PHILIP ET AL) 15 September 1998 (1998-09-15) figures 4,7	1,22,35, 42,43 2-21, 23-34, 36-41
P,X	WO 00 02410 A (DEML REINHARD ;SIEMENS AG (DE); WAHLER JOSEF (DE)) 13 January 2000 (2000-01-13) figure 2	1,22,35, 42,43
X A	WO 94 03004 A (ITALTEL SPA ;CANATO LUIGI (IT); GALLASSI GIORGIO (IT); MORGANTI MI) 3 February 1994 (1994-02-03) figure 2	1 2-43
	--- -/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

26 June 2000

Date of mailing of the international search report

04/07/2000

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3018

Authorized officer

Meurisse, W

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 00/06481

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>PERES M: "SCHALTEN WIE IN DER FORMEL 1" ELEKTRONIK, DE, FRANZIS VERLAG GMBH. MUNCHEN, vol. 47, no. 14, July 1998 (1998-07), pages 68-78, XP000669736 ISSN: 0013-5658 page 74, left-hand column, paragraph 4 -right-hand column, paragraph 4 page 78, left-hand column, paragraph 4 -right-hand column, paragraph 3 -----</p>	1-43
A	<p>US 5 781 320 A (BYERS CHARLES CALVIN) 14 July 1998 (1998-07-14) figure 3 -----</p>	6, 22, 35, 42

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/06481

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 5809022	A	15-09-1998	EP 0797373 A JP 10065742 A	24-09-1997 06-03-1998
WO 0002410	A	13-01-2000	DE 19829822 A	05-01-2000
WO 9403004	A	03-02-1994	IT 1259036 B DE 69309471 D DE 69309471 T EP 0653132 A	11-03-1996 07-05-1997 06-11-1997 17-05-1995
US 5781320	A	14-07-1998	NONE	

This Page Blank (uspto)